Hintikka Games for PCTL on Labeled Markov Chains

Harald Fecher
Informatik, Universität Freiburg
fecher@informatik.uni-freiburg.de

Michael Huth Nir Piterman Daniel Wagner
Dept. of Computing, Imperial College London
{mrh, nir.piterman, dwagner}@doc.ic.ac.uk

Abstract

We present Hintikka games for formulae of the probabilistic temporal logic PCTL and countable labeled Markov chains as models, giving an operational account of the denotational semantics of PCTL on such models. Winning strategies have a decent degree of compositionality in the parse tree of a PCTL formula and express the precise evidence for truth or falsity of a PCTL formula. We also prove the existence of monotone winning strategies that are almost finitely representable. Thus this work serves as a foundation for witness and counterexample generation in probabilistic model checking and for a uniform treatment of abstraction-based probabilistic model checking through games.

This work is also of independent interest as it displays a subtle interplay between Büchi acceptance conditions on infinite plays, the strictness or non-strictness of probability thresholds in Strong and Weak Until PCTL formulae in "GreaterThan" normal form, and a finite-state approximation lemma for Strong Until formulae with strict thresholds.

1 Introduction

Motivation. Countable labeled Markov chains [14, 6] are an important class of stochastic processes for the modeling of probabilistic systems. PCTL [12] is a probabilistic temporal logic whose formulae ϕ can express practically relevant specifications, e.g. "with probability at least $1-1/100$, a device will be elected leader" may be a requirement within a telecommunications standard such as [1], and can be written as $[\mathsf{tt}\,\mathsf{U}\,\mathsf{someLeaderElected}]_{\geq 1-1/100}$ in PCTL. A denotational semantics $[\![\phi]\!]_M$ over labeled Markov chains M then renders truth or falsity of ϕ, where $[\![\phi]\!]_M$ is the set of states in M at which ϕ is true.

Algorithms that compute this truth value require sufficient information about the probabilities of sets of paths that satisfy path sub-formulae of ϕ such as "atomic property q is true until atomic property r becomes true" (written $q\,\mathsf{U}\,r$ in linear-time temporal logic). PCTL formulae attain Boolean truth values by casting such path probabilities through the use of thresholds, e.g. "with probability at least $1/2$, atomic property q remains to be true until atomic property r becomes true". This casting also allows the approximation of the probabilities of path sets: if an incremental computation of the above probability attains a value $\geq 1/2$, the computation may safely and conclusively terminate.

A source of complexity in probabilistic model checking of PCTL on finite-state labeled Markov chains is that the familiar fixed-point characterization of Until formulae

$$q\,\mathsf{U}\,r \equiv r \vee (q \wedge \mathsf{X}\,(q\,\mathsf{U}\,r)) \tag{1}$$

– saying that "either r is true at present; or q is true at present and the Until formula is true at the next state" – cannot be carried through the casting process of thresholds in PCTL in that same simple manner, since probabilistic dependencies may not reflect this compositional interpretation of Boolean connectives. This impediment to efficient model-checking algorithms poses also an obstacle to the synthesis of meaningful and compact evidence for the truth or falsity of $s \in [\![\phi]\!]_M$. This is in striking contrast to the situation for the linear-time temporal logic LTL, where such evidence may be given by a finite state path, possibly followed by a finite state path that loops [8, Chap. 9]. This is furthermore in contrast to the situation for branching-time temporal logics such as CTL and the μ-calculus, where such evidence is attainable by a characterization of $s \in [\![\phi]\!]_M$ through 2-person games that are determined (i.e. won by exactly one of the two players), and through the synthesis of a winning strategy for the resulting game [9].

Finitary evidence. The benefit of winning strategies as complete and rigorous evidence of truth or falsity is appealing but winning strategies, as objects, may consume too much space or their synthesis may require too much time. One way in which to address this is to approximate a winning strategy with a compact object that still retains some but not necessarily all evidence of truth or falsity. Alternatively, one may seek abstract representations of the winning strategy that won't lose any precision in terms of the evidence of truth or falsity captured by the winning strategy.

Conservative abstraction techniques (e.g., [18] and [8,

Chap. 13]) appear to fall into the latter category. If state a in model A abstracts state c in model M such that $a \in [\![\phi]\!]_A$ implies $c \in [\![\phi]\!]_M$, then a winning strategy for a game that captures $a \in [\![\phi]\!]_A$ may serve as complete evidence for the truth of $c \in [\![\phi]\!]_M$. In particular, if A is finite-state then the game and any winning strategies for $a \in [\![\phi]\!]_A$ may be representable as finitary objects. Conversely, if winning strategies in a game for $c \in [\![\phi]\!]_M$ have finitary representation even for infinite-state models M, such a representation may be interpretable as a finite-state model A that abstracts M and satisfies ϕ.

In that context, it is of interest that for existentially quantified formulae of CTL such finite-state abstractions won't exist in general for abstractions that are like 3-valued Kripke structures [4]. It is a routine observation that the examples of "incompleteness" of abstraction in loc. cit. carry over to the world of labeled Markov chains and 3-valued, finite-state, labeled Markov chains as abstractions as soon as the abstraction relation preserves the positivity of path probabilities, e.g. as is the case for the probabilistic simulation of Larsen & Skou [19].

The work reported here means to establish firm foundations on which questions about the existence and computation of finitary evidence of truth of PCTL formulae, questions about the existence of finite-state abstractions that witness such truth, and questions about connections between witnessing abstractions and winning strategies can be phrased and studied.

Hintikka games. We now sketch the idea behind Hintikka games [13]. The semantics of first-order logic over models is defined as a Tarskian notion of truth: $\models$ is a formally defined predicate between models and formulae of first-order logic and "property ϕ is true in model M" is defined as "$M \models \phi$ holds". For each model M, a Hintikka game $\mathsf{G}(M, \phi)$ involves two players, Verifier (who wants to prove that M satisfies ϕ) and Refuter (who wants to prove that M does not satisfy ϕ). For example, in game $\mathsf{G}(M, \phi_1 \wedge \phi_2)$ Refuter has initial control and chooses a move to the continuation game $\mathsf{G}(M, \phi_1)$ or a move to the continuation game $\mathsf{G}(M, \phi_2)$. So Refuter is a "universal" player. Dually, in game $\mathsf{G}(M, \exists x\, \phi)$, Verifier is in initial control, chooses an element a of the structure in M, binds x to a, and moves to the continuation game $\mathsf{G}(M[x \mapsto a], \phi)$ for the model that is M but with x interpreted as a. So Verifier is an "existential" player. Sequences of such moves generate plays which are always finite since the continuation games involve proper subformulae. Eventually, a game of form $\mathsf{G}(M, R(t_1, \ldots, t_2))$ is reached for $n$ary relation symbol R and terms t_i. Verifier wins that game if the interpretation of the tuple of terms in M is contained in the interpretation of relation R in M; otherwise Refuter wins.

Strategies for both players are objects that allow them to make necessary choices for determining continuation games. For example, Verifier needs to make choices at disjunctions and existential quantifiers. A strategy σ is winning for a player if all plays played according to the choices offered by strategy σ are won by that player. Since all plays for first-order logic are finite, classical game theory guarantees that games $\mathsf{G}(M, \phi)$ are determined: exactly one of the two players has a winning strategy for that game. It is well known that in ordinary set theory ZF the assumption of the Axiom of Choice is equivalent to that

> **(Correspondence)** "Verifier wins the game $\mathsf{G}(M, \phi)$" if, and only if, "$M \models \phi$ holds".

holds. So one gets an operational and "small-step" account of truth in first-order logic from the Axiom of Choice.

In this paper we also rely on the Axiom of Choice for the composition of winning strategies for Until formulae with non-strict probability thresholds from countably many winning strategies for Until formulae with strict probability thresholds in proving **(Correspondence)** in our setting of PCTL and countable labeled Markov chains. This dependency appears to vanish for finite-state models and for PCTL formulae whose thresholds are never in control of the universal Refuter. The latter is of interest since *any* PCTL formula can be rewritten with the help of small perturbations of thresholds that won't diminish their practical value to specifiers but avoid the need for universally interpreted probability thresholds. Our games retain to idea of Verifier and Refuter as being existential and universal players (respectively), and of both having to make choices of either sub-formulae or of structural elements, which for PCTL turn out to be *sub-distributions* that approximate transition distributions in labeled Markov chains.

Contributions of paper. We formulate 2-person Hintikka games between a Verifier and a Refuter for state s and PCTL formula ϕ in a countable, labeled Markov chain M. We prove that these games are won by Verifier if, and only if, state s satisfies ϕ in M; and won by Refuter if, and only if, state s does not satisfy ϕ in M. In particular these games are determined. We then show that winning strategies can be assumed to have structural properties that make them amenable to finitary representations. We also show that such finitary representations have resemblance to *finite-state* abstract models that witness truth of a PCTL formula on the model they abstract.

Outline of paper. In Section 2 we discuss related work. In Section 3 we review the familiar denotational semantics of PCTL for countable labeled Markov chains as models, and prove a finite-state approximation lemma for (strong) Until formulae with non-strict thresholds under that semantics. In Section 4 the game semantics for PCTL over countable labeled Markov chains is being defined and these games are shown to be determined and to capture precisely the denotational semantics of PCTL. In Section 5 we discuss what

structural properties one may assume in winning strategies for our games. A discussion of the relevance of our results to finding finite abstractions is contained in Section 6, and we conclude in Section 7.

2 Related work

Model checking formulae of the μ-calculus [17] has an efficient reduction to determining the winning regions in parity games [9], 2-person games with a parity acceptance condition for infinite plays. Parity games are formulated *independently* from the notion of model and semantics of the μ-calculus, but winning regions of parity games are expressible through model checks of μ-calculus formulae. Our Hintikka games, in contrast, have Büchi acceptance conditions and are formulated in terms of labeled Markov chains and the familiar PCTL semantics.

In [10] a quantitative μ-calculus with an explicit discount operator, and with models whose transitions are labeled with discount factors has non-negative real numbers as results of model checks. Quantitative parity games are developed and shown to correspond to model checks for formulae of the quantitative μ-calculus in the same manner as for their qualitative variants above. However, winning strategies are no longer memoryless in general as they may have to "make up" for discount factors encountered en-route in a play – even in games with finite set of configurations.

Another quantitative μ-calculus (qMμ) is studied in [21], where models contain both non-deterministic and probabilistic choice but no discounting. A denotational semantics generalizes Kozen's familiar semantics to that logic. For any finite-state model and formula of qMμ the authors propose a probabilistic analogue of parity games, show that this game is determined, prove that its game value equals that of the denotational semantics for the model and formula in question, and establish that there exist memoryless winning strategies in this game.

The interplay of probabilistic approximation and logic, originally investigated in [15, 16], is being explored for continuous-state labeled Markov processes in [7]. Each model is being approximated by a chain of finite-state models ordered by a transitive simulation relation. The logic contains a constant for truth, conjunction, and the equivalent of the PCTL operator $\mathsf{X}[\cdot]_{>p}$ in the process setting. This approximating chain is linked to the logic as follows: the model satisfies a formula if, and only if, there is a model in its approximating chain that satisfies it. The chain is being built from two parameters (as in our Lemma 1): one for the temporal depth of the unfolding of the model, another for the precision with which transition probabilities of the model are represented in that truncated approximation. Thus, if a continuous-state model satisfies a formula, there is a finite-state simulation of that model that witnesses this.

This ability of witnessing truth in finite-state abstractions seems to disappear for branching-time and probabilistic logics as soon as certain fixed-points are added to the logic. In [4] one finds an infinite-state Kripke structure that satisfies an *existentially* quantified Weak Until formula but for which no simulating 3-valued variant of Kripke structure satisfies that Weak Until formula, for any sensible notion of simulation: there is no finite-state simulation of that infinite-state Kripke structure that satisfies this formula. It is easily seen that these results carry over into the world of labeled Markov chains if simulations preserve positive probabilities of path sets, as is e.g. the case for the probabilistic simulation of Larsen & Skou [19]. Note that in the setting of labeled Markov chains the existential and universal path quantifier collapse into a single probabilistic one.

In [11] finite-state (discrete-time) labeled Markov chains and probabilistic CTL (PCTL) are considered in their standard semantics, and different forms of evidence being developed for documenting the falsity of a PCTL formula in a given state. One form computes those paths that contribute most to the falsity of a formula. Another form computes most probable subtrees to gain more precise diagnostic evidence. Both forms, studied for Strong and Weak Until, are supported with shortest-path type algorithms for computing such evidence. In [2] the line of work from [11] is being pushed into the world of Markov decision processes, with a focus on upwards-bounded probability thresholds in PCTL formulae – whereas we study the downwards-bounded case without loss of generality. The shortest-path algorithms in [2] are then combined with AND/OR trees in order to filter the computed set of paths to one with high explanatory value, and to compute the probability of that filtered path set. We believe that our Hintikka games provide a suitable foundation for understanding the trade-off between the precision of extant and future forms of evidence and the complexity of their supporting algorithms.

Games are an attractive candidate for describing model checks, abstraction between models, and model checks on abstract models within a single formalism through satisfaction games, refinement games, and abstract satisfaction games, respectively. This has been developed for Kripke structures, focused transition systems as their abstractions, and the modal mu-calculus in [4]. Tree automata are equally attractive in that regard, as demonstrated in [5]. It would therefore be of interest to generalize our Hintikka games so that they are formulated independently of a notion of model and logic.

Stochastic games [3] have not only adversarial players, e.g. the two players 0 and 1 for parity games, but an additional random player $.5$ whose game moves are determined by probability distributions. For such so-called 2.5-player games one can consider the usual acceptance conditions for infinite plays. For reachability, 2.5-player stochastic

games have successfully been applied in improving the precision of abstractions of Markov decision processes [18]: one adversarial player controls the partition of the concrete state space, the other adversarial player controls the nondeterministic choices inherent in the concrete Markov decision process. It would be of interest to see if 2.5-player stochastic games with Büchi acceptance conditions provide finite-state abstractions that can witness the truth of PCTL formulae of labeled Markov chains they abstract.

In [22], probabilistic bisimulation is modified to a notion of bisimulation "up to ϵ" that is shown to be of use in quantitative analysis of confinement problems in security. This notion is no longer transitive and bisimulation "up to 0" does not coincide with probabilistic bisimulation for infinite-state systems. It would be of interest to understand whether the ϵ-moves in our Hintikka games correspond to approximative versions of probabilistic simulations.

3 Preliminaries

(Countable) Labeled Markov chains M over a set of atomic propositions $\mathbb{AP}$ are triples (S, P, L), where S is a countable set of states, $P\colon S \times S \to [0,1]$ is a countable stochastic matrix such that the countable sum of nonnegative reals $\sum_{s' \in S} P(s,s')$ converges to 1 for all $s \in S$, and $L\colon \mathbb{AP} \to \mathbb{P}(S)$ is a labeling function where $L(q)$ is the set of states at which atomic proposition q is true. We say that M is finitely branching iff for all $s \in S$ the set $\{s' \in S \mid P(s,s') > 0\}$ is finite. A path π from state s in M is an infinite sequence of states $s_0 s_1 \ldots$ with $s_0 = s$ and $P(s_i, s_{i+1}) > 0$ for all $i \geq 0$. For $Y \subseteq S$, we write $P(s,Y)$ as a shorthand for the (possibly infinite but well defined) sum $\sum_{s' \in Y} P(s,s')$.

The syntax of PCTL is given in Fig. 1. Path formulae α are wrapping PCTL formulae into "LTL" operators for Next, (strong) Until, and Weak Until. Path formulae are interpreted as predicates over paths of M. The semantics is defined as usual: a path $\pi = s_0 s_1 \ldots$ satisfies

- $\mathsf{X}\,\phi$ iff $s_1 \in [\![\phi]\!]_M$
- $\phi\,\mathsf{U}^{\leq k}\psi$ iff there is a $l \in \mathbb{N}$ such that $l \leq k$, $s_l \in [\![\psi]\!]_M$ and for all $0 \leq j < l$ we have $s_j \in [\![\phi]\!]_M$
- $\phi\,\mathsf{W}^{\leq k}\psi$ iff for all $l \in \mathbb{N}$ such that $0 \leq l \leq k$ we have either $s_l \in [\![\phi]\!]_M$ or there is $0 \leq j \leq l$ with $s_j \in [\![\psi]\!]_M$

Until formulae $\phi\,\mathsf{U}^{\leq k}\psi$ are *strong* untils since paths that satisfy such a formula have to maintain temporary invariant ϕ until they reach a state satisfying ψ, and such a state has to be reached within finite transitions, and also within k transitions if $k \neq \infty$. Weak Until formulae $\phi\,\mathsf{W}^{\leq k}\psi$ are *weak* untils since reaching a state satisfying ψ is optional if ϕ is an invariant on the path $s_0 s_1 \ldots s_k$, which is understood to be π when $k = \infty$. We record the familiar duality between

ϕ, ψ ::=	*PCTL formulae*	α ::=	*Path formulae*
q	Atom	$\mathsf{X}\,\phi$	Next
$\neg\phi$	Negation	$\phi\,\mathsf{U}^{\leq k}\psi$	Until
$\phi \wedge \psi$	Conjunction	$\phi\,\mathsf{W}^{\leq k}\psi$	Weak Until
$[\alpha]_{\bowtie p}$	Path Probability		

Figure 1. Syntax of PCTL, where $\mathsf{q} \in \mathbb{AP}$, $k \in \mathbb{N} \cup \{\infty\}$, $p \in [0,1]$, and $\bowtie \in \{<, \leq, >, \geq\}$

(strong) Until and Weak Until:

$$\neg(\phi\,\mathsf{U}\,\psi) \equiv (\neg\psi)\,\mathsf{W}\,(\neg\phi \wedge \neg\psi) \tag{2}$$

PCTL formulae wrap path formulae with probability thresholds (turning predicates on paths into predicates on states), and may add a propositional logic layer on top of that, which may then be used to build up new Path formulae. We write $\phi\,\mathsf{U}\,\psi$ as a shorthand for $\phi\,\mathsf{U}^{\leq\infty}\psi$, and $\phi\,\mathsf{W}\,\psi$ as shorthand for $\phi\,\mathsf{W}^{\leq\infty}\psi$. The operators $\vee$ (disjunction) and $\rightarrow$ (implication) are derived as usual. Let ff be an abbreviation for any $[\alpha]_{>1}$, and tt a shorthand for any $[\alpha]_{\geq 0}$. For labeled Markov chain $M = (S, P, L)$, the denotational semantics of PCTL formula ϕ is a subset $[\![\phi]\!]_M$ of S. We write $[\![\phi]\!]$ if M is clear from the context and define $[\![\phi]\!]$ by structural induction, as usual:

$$[\![q]\!] = L(q) \qquad [\![\phi \wedge \psi]\!] = [\![\phi]\!] \cap [\![\psi]\!]$$
$$[\![\neg\phi]\!] = S \setminus [\![\phi]\!] \qquad [\![[\alpha]_{\bowtie p}]\!] = \{s \in S \mid \mathsf{Prob}_M(s,\alpha) \bowtie p\}$$

where $\mathsf{Prob}_M(s,\alpha)$ is the probability of the measurable set $\mathsf{Path}(s,\alpha)$ of paths in M that begin in s and satisfy the path formula α. Note that the semantics of PCTL and Path formulae is mutually recursive, reflecting the mutual recursion of their syntax. We say that PCTL formulae ϕ and ψ are semantically equivalent iff for all labeled Markov chains M we have $[\![\phi]\!]_M = [\![\psi]\!]_M$.

Example 1 *For the labeled Markov chain M in Figure 2 we have $[\![[q\,\mathsf{U}\,r]_{\geq 1/2}]\!]_M = \{s_0, s_1\}$ and for the labeled Markov chain $M_2^{s_0}$ in Figure 3 we have $[\![[q\,\mathsf{W}\,r]_{\geq 5/9}]\!]_M = \{s_0, s_0 s_1, s_0 s_1 s_1, s_0 s_0, s_0 s_0 s_1, s_0 s_0 s_0\}$.*

Each PCTL formula ϕ is semantically equivalent to a PCTL formula in "GreaterThan" normal form obtained by replacing all occurrences of the form $[\alpha]_{<p}$ in ϕ with the PCTL formula $\neg[\alpha]_{\geq p}$, and by replacing any occurrences of the form $[\alpha]_{\leq p}$ in ϕ with the PCTL formula $\neg[\alpha]_{>p}$. For example, the "GreaterThan" normal form of the formula $[[\mathsf{X}\,[q\,\mathsf{U}\,r]_{<1/3}]_{\leq 1/2}\,\mathsf{U}\,r]_{>1/4}$ is $[\neg[\mathsf{X}\,\neg[q\,\mathsf{U}\,r]_{\geq 1/3}]_{>1/2}\,\mathsf{U}\,r]_{>1/4}$.

Assumption 1 (GreaterThan) *Without loss of generality, PCTL of Fig. 1 is restricted to "GreaterThan" normal form, i.e., $\bowtie \in \{\geq, >\}$.*

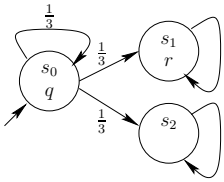

Figure 2. Labeled Markov chain M with $s_0 \in [\![q \mathsf{U} r]\!]_{\geq 1/2}]\!]_M$, since $\mathsf{Prob}_M(s_0, q \mathsf{U} r) = 1/2$.

We now state and prove a finite-state approximation lemma for the validity of Until formulae with non-strict probability thresholds at states of labeled Markov chains. This lemma will be crucial in proving that our game semantics of PCTL, developed in Section 4, captures exactly the denotational semantics we defined above.

Definition 1 (Finite Unfoldings) *Let $M = (S, P, L)$ be a labeled Markov chain. For each $i \in \mathbb{N}$ and $s_0 \in S$ define the labeled Markov chain $M_i^{s_0} = (S_i, P_i, L_i)$, a random tree with root s_0, as follows: unfold M from s_0 as a full tree of depth i, where edges have positive probability according to P. This may duplicate states but such duplicates will satisfy the same atomic propositions. States at level i have a self-loop with probability 1. The probability measures $P(s, \cdot)$ at levels $< i$ are those in M. For each $j \in \mathbb{N}$ we restrict $M_i^{s_0}$ to the finite-branching, and so finite-state, labeled Markov chain $M_{i,j}^{s_0} = (S_{i,j}, P_{i,j}, L_{i,j})$ with one additional state t_{sink} which only satisfies $\mathtt{tt}$ but no other $q \in \mathbb{AP}$: for each $s \in S_i$, let $t_1, t_2, \ldots$ be an enumeration of $\{t_k \in S_i \mid P(s, t_k) > 0\}$ such that $P(s, t_k) \geq P(s, t_{k+1})$ for all $k \in \mathbb{N}$, then $P_{i,j}$ is obtained from P_i by setting $P_{i,j}(s, t_k) = P_i(s, t_k)$ for $k \leq j$, $P_{i,j}(s, t_{\text{sink}}) = 1 - \Sigma_{k=1}^{j} P_{i,j}(s, t_k)$ and $P_{i,j}(t_{\text{sink}}, t_{\text{sink}}) = 1$; state set $S_{i,j}$ consists of those s reachable from s_0 via $P_{i,j}$, and $L_{i,j}$ is L_i restricted to set $S_{i,j}$ and extended to the new state t_{sink}.*

Example 2 *The unfolding $M_2^{s_0}$ for the labeled Markov chain M of Figure 2 is depicted in Figure 3.*

Lemma 1 (Finite-State Approximation) *Let $M = (S, P, L)$ be a labeled Markov chain, $q, r \in \mathbb{AP}$, and $p \in [0, 1]$. Then $s \in [\![q \mathsf{U} r]\!]_{\geq p}]\!]_M$ iff for all $n \in \mathbb{N}$ there are $k, l \in \mathbb{N}$ with $s \in [\![q \mathsf{U} r]\!]_{>p-1/n}]\!]_{M_{k,l}^s}$.*

Example 3 *Consider the labeled Markov chain in Fig. 2. Probability $\mathsf{Prob}_M(s_0, q \mathsf{U} r) = 1/2$ is attained by paths of increasing length, as the value of the infinite sum $\Sigma_{j=1}^{\infty}(1/3)^j$. However, for every $n \in \mathbb{N}$ there exists some $i \in \mathbb{N}$ such that $\Sigma_{j=1}^{i}(1/3)^j > 1/2 - 1/n$ and where that finite sum is attainable in a finite unfolding of M. For example, for $M_2^{s_0}$ in Fig. 3 the probability of $q \mathsf{U} r$ at s_0 is $\frac{4}{9}$ so for every $n < 18$ we have $s_0 \in [\![q \mathsf{U} r]\!]_{>1/2-1/n}]\!]_{M_2^{s_0}}$. In $M_4^{s_0}$ the probability of $q \mathsf{U} r$ at s_0 is $\frac{13}{27}$ and so for every $n < 54$ we have $s_0 \in [\![p \mathsf{U} q]\!]_{>1/2-1/n}]\!]_{M_4^{s_0}}$. Lemma 1 promises that for every (countable) labeled Markov chain there is a similar approximation.*

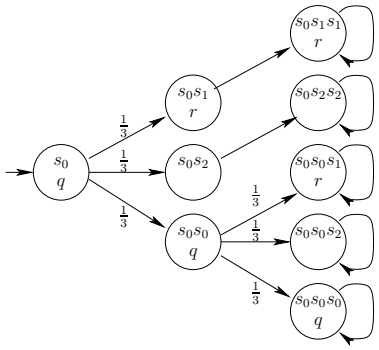

Figure 3. Unfolding $M_2^{s_0}$ of the labeled Markov chain of Figure 2 up to depth two.

Lemma 1 has a dual version, required in the proof of Theorem 1: for labeled Markov chain $M = (S, P, L)$, $q, r \in \mathbb{AP}$, and $p \in [0, 1]$: $s \notin [\![q \mathsf{W} r]\!]_{>p}]\!]_M$ iff for all $n \in \mathbb{N}$ there are $k, l \in \mathbb{N}$ with $s \notin [\![q \mathsf{W} r]\!]_{\geq p+1/n}]\!]_{M_{k,l}^s}$.

4 Game semantics

Let $M = (S, P, L)$ be a labeled Markov chain over set of atomic propositions $\mathbb{AP}$. For each state $s \in S$ and PCTL formula ϕ we define a 2-person Hintikka game $\mathsf{G}_M(s, \phi)$. As already mentioned, these games are played between two players $\mathtt{V}$ (the Verifier) and $\mathtt{R}$ (the Refuter). After having defined these games and their winning conditions, we show that each game $\mathsf{G}_M(s, \phi)$ is won by player $\mathtt{V}$ iff $s \in [\![\phi]\!]_M$; and won by player $\mathtt{R}$ iff $s \notin [\![\phi]\!]_M$. In particular, each game $\mathsf{G}_M(s, \phi)$ is *determined*, exactly one of the players $\mathtt{V}$ and $\mathtt{R}$ wins that game.

The game $\mathsf{G}_M(s, \phi)$ has as set of configurations

$$\mathsf{Cf}_M(s, \phi) = \{\langle s', \psi, \mathtt{C}\rangle \mid s' \in S,\ \psi \in \mathsf{cl}(\phi),\ \mathtt{C} \in \{\mathtt{R}, \mathtt{V}\}\}$$

where we define the set of PCTL formulae $\mathsf{cl}(\phi)$, the *closure of ϕ*, below. There is a distinguished initial configuration $\langle s, \phi, \mathtt{V}\rangle \in \mathsf{Cf}_M(s, \phi)$. Plays in game $\mathsf{G}_M(s, \phi)$ are finite or infinite sequences of elements in $\mathsf{Cf}_M(s, \phi)$ starting in the initial configuration $\langle s, \phi, \mathtt{V}\rangle$. A play is generated by game moves, specified in detail below. Their intuition is similar to that in Hintikka games for first-order logic, as described in the introduction.

Our game semantics treats Boolean connectives in the same manner as Hintikka games for first-order logic (here we take the point of view of Verifier): proving truth of formula ϕ at state s amounts to winning the game from configuration $\langle s, \phi, \mathtt{V}\rangle$. In order to prove a conjunction we allow Refuter to choose which branch of the conjunction to prove. In order to handle negation, the game continues in the same state but with the unnegated formula and a swap of the role of players, thus attempting to show that Refuter cannot win from the unnegated formula.

In games for branching-time logics such as CTL or the μ-calculus (see e.g. [23]), the universal quantification in $\forall\mathsf{X}$ is resolved by Refuter's choice of a successor state; and the existential quantification in $\exists\mathsf{X}$ is resolved by Verifier supplying one successor state, both as familiar from the case of quantifiers in first-order logic. For PCTL, however, things are more complicated. The next operator $[\mathsf{X}\,\phi]_{\bowtie p}$ includes a promised probability $\bowtie p$, "at least p" or "more than p". Verifier now resolves this "probabilistic quantification" by showing how to re-distribute the required probability between the successors of the state.

In *qualitative* games, until operators are resolved by using the logical equivalence in (1) – and similarly for weak until operators. The only problem in adopting this for PCTL is in the possibility of deferring promises forever. For games in qualitative settings this is typically handled by fairness, but for PCTL fairness is not strong enough:

Example 4 *PCTL formula $[q\,\mathsf{U}\,r]_{\geq 1/2}$ holds at state s_0 in the labeled Markov chain shown in Figure 2. But in order to prove this we have to appeal to the entire infinite sum $\Sigma_{i=1}^{\infty}(1/3)^i$. Any fairness constraint forcing a transition from s_0 into $\{s_1, s_2\}$ would cut that infinite sum down to a finite one, failing to prove that formula for state s_0.*

However, allowing to defer the satisfaction of the strong until indefinitely is unsound. For any $\bowtie p$ being $> p$ with $p \in [0.5, 1]$, and any $\bowtie p$ being $\geq p$ with $p \in (0.5, 1]$, the PCTL formula $[q\,\mathsf{U}\,r]_{\bowtie p}$ does not hold at s_0 but allowing Verifier to delay promises forever may be unsound. For $\bowtie p$ being > 0.5, e.g., Verifier could supply the promise $1/3$ immediately, promising more than $1/6$ in the future, and then – by deferring the promise indefinitely – Verifier could win game $\mathsf{G}_M(s_0, [q\,\mathsf{U}\,r]_{>0.5})$.

To address this problem we add a special ϵ-move as well as acceptance conditions for infinite plays. If the probability is at least p, player V should be able to prove that it is greater than $p - \epsilon$ for every $\epsilon > 0$. On the other hand, if the probability is strictly less than p then there exists an ϵ for which it is at most $p - \epsilon$. Thus, player R chooses the ϵ and player V proves in finite time (appealing to Lemma 1) that she can get as close as needed to the bound. The same intuition (but dual) works for *Weak* Until, when the Weak Until formula in question does *not* hold.

We next define the notion of the closure $\mathsf{cl}(\phi)$ of a PCTL formula ϕ, which is the union of two sets of PCTL formulae. The first set $\mathsf{cl}_1(\phi)$ is the actual set of sub-PCTL-formulae of ϕ, including ϕ itself. The second set $\mathsf{cl}_2(\phi)$ consists of all formulae $[\alpha]_{\bowtie p'}$ such that either

(a) α is $\psi_1\,\mathsf{U}\,\psi_2$, $\bowtie$ is $>$, and for some $p \in [0,1]$ and $\bowtie' \in \{>, \geq\}$ we have $[\alpha]_{\bowtie' p} \in \mathsf{cl}_1(\phi)$,
(b) α is $\psi_1\,\mathsf{W}\,\psi_2$, $\bowtie$ is $\geq$, and for some $p \in [0,1]$ and $\bowtie' \in \{>, \geq\}$ we have $[\alpha]_{\bowtie' p} \in \mathsf{cl}_1(\phi)$,
(c) α is $\psi_1\,\mathsf{U}^{\leq k'}\psi_2$ and for some $p \in [0,1]$ and a finite $k > k'$ we have $[\psi_1\,\mathsf{U}^{\geq k}\psi_2]_{\bowtie p} \in \mathsf{cl}_1(\phi)$,
(d) α is $\psi_1\,\mathsf{W}^{\leq k'}\psi_2$ and for some $p \in [0,1]$ and a finite $k > k'$ we have $[\psi_1\,\mathsf{U}^{\geq k}\psi_2]_{\bowtie p} \in \mathsf{cl}_1(\phi)$

The second set $\mathsf{cl}_2(\phi)$ allows us to replace any probability thresholds p with other values $p' \in [0,1]$ and finite time bounds with smaller ones, but to allow this in such a manner that it is consistent with the above intuition behind ϵ-moves:

- (strong) Until formulae with non-strict bounds may change to (strong) Until formulae with strict bounds
- Weak Until formulae with strict bounds may change to Weak Until formulae with non-strict bounds, and
- the finite time bounds in bounded untils should be allowed to decrease.

The difference between the strong and weak untils stems from their duality, the negation of a Weak Until formula is a (strong) Until formula and vice versa. Thus, a Weak Until formula with strict bound is the negation of a (strong) Until formula with non-strict bound. When Refuter is trying to disprove a Weak Until formula with strict bound, she is in fact trying to prove the dual (strong) Until formula with non-strict bound, and requires the same possible moves for the non-strict bound and strict bound versions.

Example 5 *Consider the following formula:*

$$\phi = [(r \wedge [\mathsf{X}\,[(p \wedge \neg r)\,\mathsf{W}\,(q \wedge \neg r)]_{\geq 1}]_{>0})\,\mathsf{W}\,\mathsf{ff}]_{>0} \quad (3)$$

Intuitively, ϕ says that there is an infinite path labeled by r such that every state on this path has a successor for which $p\,\mathsf{W}\,q$ holds on (almost) all paths on which r does not hold during the verification of $p\,\mathsf{W}\,q$. Let $\alpha = (p \wedge \neg r)\,\mathsf{W}\,(q \wedge \neg r)$, $\beta = \mathsf{X}\,[\alpha]_{>0}$, and $\gamma = (r \wedge [\beta]_{>0})\,\mathsf{W}\,\mathsf{ff}$. The closure of ϕ is:

$$\mathsf{cl}(\phi) = \left\{ \begin{array}{l} \phi, [\gamma]_{\geq b}, \mathsf{ff}, (r \wedge [\beta]_{>0}), \\ [\beta]_{>0}, [\alpha]_{\geq b}, (p \wedge \neg r), \\ p, \neg r, r, (q \wedge \neg r), q \end{array} \middle| \; b \in [0,1] \right\}$$

As γ appears in ϕ with a strict bound, it is in the closure of ϕ with its original bound as well as with all possible non-strict bounds. As α appears in ϕ with a non-strict bound, it appears in the closure of ϕ only with non-strict bounds.

Similarly, for formula $\phi = [q \,\mathsf{U}\, r]_{\geq 1/2}$ *we have* $\mathsf{cl}(\phi) = \{\phi, q, r, [q \,\mathsf{U}\, r]_{>b} \mid b \in [0,1]\}$. *As* ϕ *is a strong until with non-strict bound it is part of* $\mathsf{cl}_1(\phi)$ *and for every possible bound* b *its strict counterpart* $[q \,\mathsf{U}\, r]_{>b}$ *is in* $\mathsf{cl}_2(\phi)$.

Subsequently, we write $!\mathtt{C}$ for the player other than $\mathtt{C}$, i.e. $!\mathtt{V} = \mathtt{R}$ and $!\mathtt{R} = \mathtt{V}$. The possible moves of game $\mathsf{G}_M(s_0, \phi)$ are defined through the moves of games $\mathsf{G}_M(s, \psi)$ by structural induction on $\psi \in \mathsf{cl}(\phi)$, simultaneously for all $s \in S$.

M1. At configurations $\langle s, [\alpha]_{>1}, \mathtt{C}\rangle$, player $!\mathtt{C}$ wins

M2. At configurations $\langle s, [\alpha]_{\geq 0}, \mathtt{C}\rangle$, player $\mathtt{C}$ wins

We may therefore assume that in subsequent moves configurations of the form $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}\rangle$ never satisfy that $\bowtie\, p$ equals ≥ 0 or > 1.

M3. At configurations $\langle s, \mathsf{q}, \mathtt{C}\rangle$:

- player $\mathtt{C}$ wins if $s \in L(\mathsf{q})$
- player $!C$ wins if $s \notin L(\mathsf{q})$

M4. At configuration $\langle s, \neg\psi, \mathtt{C}\rangle$, the next configuration is $\langle s, \psi, !\mathtt{C}\rangle$

So move M4 removes the negation from the formula but also swaps the role of players.

M5. At configuration $\langle s, \psi_1 \wedge \psi_2, \mathtt{C}\rangle$, player $!\mathtt{C}$ can choose as next configuration either $\langle s, \psi_1, \mathtt{C}\rangle$ or $\langle s, \psi_2, \mathtt{C}\rangle$

So player $!\mathtt{C}$ chooses a conjunct and the game continues with that conjunct instead of the conjunction.

M6. At configuration $\langle s, [\,\mathsf{X}\,\psi]_{\bowtie p}, \mathtt{C}\rangle$, player $\mathtt{C}$ chooses a subset $Y \subseteq S$ satisfying $P(s, Y) \bowtie p$; then player $!\mathtt{C}$ chooses some $s' \in Y$:

- if $P(s, s') = 0$, player $!\mathtt{C}$ wins
- otherwise, $P(s, s') > 0$ and the next configuration is $\langle s', \psi, \mathtt{C}\rangle$

Move M6 is well defined. There is a non-empty set Y with $P(s, Y) \bowtie p$ as $p \in [0, 1]$, $P(s, \cdot)$ has mass one, and $\bowtie\, p$ is neither equal to > 1 nor to ≥ 0.

M7. At configuration $\langle s, [\psi_1 \,\mathsf{U}\, \psi_2]_{\geq p}, \mathtt{C}\rangle$, player $!\mathtt{C}$ chooses some $n \in \mathbb{N}$ such that $p - 1/n \geq 0$ with resulting next configuration $\langle s, [\psi_1 \,\mathsf{U}\, \psi_2]_{>p-1/n}, \mathtt{C}\rangle$

In move M7 such a choice is possible since p cannot be 0. The intuition is that $[p, 1] = \bigcap_{n\in\mathbb{N}}(p - 1/n, 1]$ so this behaves like a *universal* quantification over $n \in \mathbb{N}$.

M8. Dually, at configuration $\langle s, [\psi_1 \,\mathsf{W}\, \psi_2]_{>p}, \mathtt{C}\rangle$, now player $\mathtt{C}$ chooses $n \in \mathbb{N}$ such that $p + 1/n \leq 1$ with resulting next configuration $\langle s, [\psi_1 \,\mathsf{W}\, \psi_2]_{\geq p+1/n}, \mathtt{C}\rangle$

In move M8 such a choice is possible since $p < 1$. The intuition is that a Weak Until with a $>$ threshold is the dual of a strong until with a $\geq$ threshold (based on (2)), so it is like an *existential* quantification over $n \in \mathbb{N}$.

M9. At configuration $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}\rangle$ where either α is $\psi_1 \,\mathsf{U}\, \psi_2$ and $\bowtie$ is $>$; or α is $\psi_1 \,\mathsf{W}\, \psi_2$ and $\bowtie$ is $\geq$

- player $\mathtt{C}$ is able to move to next configuration $\langle s, \psi_2, \mathtt{C}\rangle$
- if player $\mathtt{C}$ did not move, player $!\mathtt{C}$ is able to move to next configuration $\langle s, \psi_1, \mathtt{C}\rangle$
- if neither player moved above, the play must proceed as follows:

  Player $\mathtt{C}$ chooses a sub-distribution $d\colon S \to [0, 1]$ such that

$$\sum_{s'\in S} d(s') > 0 \quad \& \quad \sum_{s'\in S} d(s') \geq p \tag{4}$$

$$\forall s' \in S\colon d(s') \leq P(s, s') \tag{5}$$

  Next, player $!\mathtt{C}$ chooses some $s' \in S$ with $d(s') > 0$ and the next configuration is $\langle s', [\alpha]_{\bowtie d(s')\cdot P(s,s')^{-1}}, \mathtt{C}\rangle$.

In move M9, sub-distribution d has positive mass, approximates the probability distribution $P(s, \cdot)$, and specifies the re-distribution of promise $\bowtie\, p$ into promised probabilities at successor states. Since $d(s') > 0$, we also have $0 < d(s') \cdot P(s, s')^{-1} \leq 1$ in move M9 by (5).

M10. At configuration $\langle s, [\alpha]_{>p}, \mathtt{C}\rangle$ where α is either $\psi_1 \,\mathsf{U}^{\leq k} \psi_2$ or $\psi_1 \,\mathsf{W}^{\leq k} \psi_2$ with $k \in \mathbb{N}$:

- if $k = 0$ and α is $\psi_1 \,\mathsf{U}^{\leq k} \psi_2$, the next configuration is $\langle s, \psi_2, \mathtt{C}\rangle$
- if $k = 0$ and α is $\psi_1 \,\mathsf{W}^{\leq k} \psi_2$, player $\mathtt{C}$ chooses as next configuration either $\langle s, \psi_1, \mathtt{C}\rangle$ or $\langle s, \psi_2, \mathtt{C}\rangle$
- if $k > 0$, the moves are defined as in M9 above; except when the last item of M9 applies, in which case the counter k in α is decreased to $k - 1$ for that next configuration $\langle s', [\alpha]_{\bowtie d(s')\cdot P(s,s')^{-1}}, \mathtt{C}\rangle$

In move M10, a Bounded Until with bound 0 has to realize ψ_2 right away; and a Bounded Weak Until with bound zero has to realize at least one of ψ_1 or ψ_2 right away.

A finite play is won as explained in M1-M10 above. In most moves, the play either ends or moves to configurations where the formula is a *proper* subformula in the closure. In a configuration with strong until with non-strict bound or weak until with strict bound the next configuration changes from non-strict to strict bound or vice versa. In a configuration with strong until with strict bound or weak until with non-strict bound the next configuration has the same path formula and threshold type, or has a proper sub-formula.

It follows that every infinite play ends with an infinite suffix of configurations that are

A1. all of the form $\langle s_i, [\psi_1 \,\mathsf{W}\, \psi_2]_{\geq p_i}, \mathtt{C}\rangle$ or
A2. all of the form $\langle s_i, [\psi_1 \,\mathsf{U}\, \psi_2]_{> p_i}, \mathtt{C}\rangle$

Configurations of these suffixes are either labeled by strong until with strict bound or weak until with non-strict bound, where the states and the exact probability bound may still change, but where neither the player C nor the subformulae ψ_1 and ψ_2 change.

Definition 2 (Acceptance conditions) *Player* V *wins all infinite plays with an infinite suffix either of type A1 above with* $\mathtt{C} = \mathtt{V}$*, or of type A2 above with* $\mathtt{C} = \mathtt{R}$*. Player* R *wins all other infinite plays: those with an infinite suffix either of type A1 when* $\mathtt{C} = \mathtt{R}$*, or of type A2 when* $\mathtt{C} = \mathtt{V}$*.*

These are Büchi type acceptance conditions, and so our games are known to be determined [20]. We use the notion of strategy for player C informally. But such strategies contain, for each configuration of a game, at most one set of choices as required by the applicable move from M1-M10.

Example 6 *Consider game* $\mathsf{G}_M(s_0, [q \,\mathsf{U}\, r]_{\geq 1/2})$*, where* M *is as in Fig. 2, and let* $\alpha = q \,\mathsf{U}\, r$*. The initial configuration is* $\langle s_0, [\alpha]_{\geq 1/2}, \mathtt{V}\rangle$*. In the first move player* R *chooses an* $n \in \mathbb{N}$ *and the next configuration is* $\langle s_0, [\alpha]_{>1/2-1/n}, \mathtt{V}\rangle$*. Then, as long as the play* $\Gamma_0\Gamma_1\ldots$ *remains in configurations of the form* $\langle s_0, [\alpha]_{>p_i}, \mathtt{V}\rangle$*, player* V *is going to choose the sub-distribution* d *with constant values* $d(s_2) = 0$ *and* $d(s_1) = \frac{1}{3} - \frac{1}{2n}$*, and dynamic value* $d(s_0) = p_i - d(s_1)$*. A simple calculation shows that as long as player* R *chooses* s_0 *as the next state (clearly, if she chooses* s_1 *she is going to lose as* $s_1 \in L(r)$*) the promised probability* $> p_i$ *is going to decrease according to the following sequence:* $p_0 = \frac{1}{2} - \frac{1}{n}$*,* $p_1 = \frac{1}{2} - \frac{3}{2n}$*,* $p_2 = \frac{1}{2} - \frac{6}{2n}$*,* $p_3 = \frac{1}{2} - \frac{15}{2n}$*, and in general* $p_i = \frac{1}{2} - \frac{3^i+3}{4n}$ *for* $i \in \mathbb{N}$*. Whenever* p_i *decreases below* $\frac{1}{3}$ *(and there is some* $i \in \mathbb{N}$ *for which this happens), player* V *still chooses* d *with* $d(s_2) = 0$ *as above but now defines* $d(s_1) = p_i$ *and* $d(s_0) = 0$*, thereby forcing player* R *to move to* s_1 *and lose. This describes a winning strategy for player* V *in game* $\mathsf{G}_M(s_0, [q \,\mathsf{U}\, r]_{\geq 1/2})$*.*

Example 7 *Although the choice of* d *in Example 6 may seem arbitrary, it meshes well with the use of Lemma 1. Consider again the game* $\mathsf{G}_M(s_0, [\alpha]_{\geq 1/2})$ *from Example 6. Suppose that in the first move player* R *chooses* $9 \in \mathbb{N}$*, and the next configuration is* $\langle s_0, [\alpha]_{>7/18}, \mathtt{V}\rangle$*. Since for the* $M^2_{s_0}$ *in Figure 3,* $\mathsf{Prob}_{M_2^{s_0}}(s_0, \alpha) = \frac{4}{9} > \frac{7}{18}$*, player* V *can use* $M_2^{s_0}$ *to guide her choices. In* $M_2^{s_0}$ *we have* $\mathsf{Prob}_{M_2^{s_0}}(s_0 s_1, \alpha) = 1$ *and* $\mathsf{Prob}_{M_2^{s_0}}(s_0 s_0, \alpha) = \frac{1}{3}$*. Player* V *uses the gap of* $\frac{1}{18}$ *and re-distributes it between the successors of* s_0*. She can choose, for example,* $d(s_1) = \frac{1}{3} - \frac{1}{54}$ *and* $d(s_0) = \frac{1}{9} - \frac{1}{54}$*. The next possible configurations are then* $\langle s_1, [\alpha]_{>17/18}, \mathtt{V}\rangle$ *and* $\langle s_0, [\alpha]_{>5/18}, \mathtt{V}\rangle$*. Player* V *identifies the resulting states with those obtained in* $M_2^{s_0}$*, here* $s_0 s_1$ *and* $s_0 s_0$ *(respectively). As* $s_0 s_1 \in [\![r]\!]_{M_2^{s_0}}$ *the first is clearly a winning configuration. From* $\langle s_0, [\alpha]_{>5/18}, \mathtt{V}\rangle$ *and the corresponding location* $s_0 s_0$ *in* $M_2^{s_0}$*, player* V *notices that* $\mathsf{Prob}_{M_2^{s_0}}(s_0 s_0 s_1, \alpha) = 1$ *and chooses* $d(s_1) = 5/18$*. The next configuration is* $\langle s_1, [\alpha]_{>15/18}, \mathtt{V}\rangle$ *(with corresponding* $s_0 s_0 s_1$ *in* $M_2^{s_0}$*) and won by supplying* r*.*

Definition 3 *1. A strategy* w *for player* C *in game* $\mathsf{G}_M(s, \phi)$ *is* winning from a configuration Γ *in that game iff player* C *wins all plays beginning in configuration* Γ *when player* C *plays according to his strategy* w *– regardless of how player* !C *plays.*

2. Player C wins *game* $\mathsf{G}_M(s, \phi)$ *iff player* C *has a strategy that is winning from configuration* $\langle s, \phi, \mathtt{V}\rangle$*.*

We can now formalize our main result that the denotational semantics of PCTL is captured exactly by the existence of winning strategies in games $\mathsf{G}_M(s, \phi)$.

Theorem 1 *Let* $M = (S, P, L)$ *be a labeled Markov chain over* $\mathbb{AP}$*,* $s \in S$*, and* ϕ *a PCTL formula. Then we have:*

1. $s \in [\![\phi]\!]_M$ *iff player* V *wins game* $\mathsf{G}_M(s, \phi)$
2. $s \notin [\![\phi]\!]_M$ *iff player* R *wins game* $\mathsf{G}_M(s, \phi)$*.*

In particular, game $\mathsf{G}_M(s, \phi)$ *is determined.*

Game $\mathsf{G}_M(s, \phi)$ is defined such that its initial configuration $\langle s, \phi, \mathtt{V}\rangle$ is owned by player V. We can define a dual game with the same moves but with initial configuration $\langle s, \phi, \mathtt{R}\rangle$. Theorem 1 and its proof then remain to be valid if we swap the role of players in both.

Example 8 *Consider game* $\mathsf{G}_M(s_0, [q \,\mathsf{U}\, r]_{>1/2})$*, where* M *is as in Fig. 2, and let* $\alpha = q \,\mathsf{U}\, r$*. From configuration* $\langle s_0, [\alpha]_{>1/2}, \mathtt{V}\rangle$*, player* V *won't move to* $\langle s_0, r, \mathtt{V}\rangle$ *as she would then lose. For the same reason, player* R *won't move to* $\langle s_0, q, \mathtt{V}\rangle$*. So if both players play strategies that are 'optimal' for them, player* V *has to choose a sub-distribution* d *at the initial configuration.*

If $d(s_2) > 0$*, player* V *loses as player* R *can then choose* s_2*. So* $d(s_2) = 0$ *for any 'optimal' strategy of player* V*. But both* $d(s_1)$ *and* $d(s_0)$ *have to be positive since otherwise the mass of* d *can be at most* $1/3$ *by (5), which would violate (4). Since player* V *plays an 'optimal' strategy,* $d(s_1) \neq 1/3$*, as otherwise player* R *could choose as next configuration* $\langle s_1, [\alpha]_{>(1/3)\cdot(1/3)^{-1}}, \mathtt{V}\rangle$ *and would then win by move M1. By (5) there is therefore* $\epsilon > 0$ *such that* $d(s_1) = 1/3 - \epsilon$*. In particular, player* R *won't choose* s_1 *as she would lose the next configuration* $\langle s_1, [\alpha]_{>1-3\epsilon}, \mathtt{V}\rangle$ *(since* $s_1 \in L(r)$*). So player* R *chooses* s_0 *and the next configuration is* $\langle s_0, [\alpha]_{>3d(s_0)}, \mathtt{V}\rangle$*. By (4),* $3d(s_0)$ *must be at least* $1/2 + 3\epsilon$ *and so player* V *promises* more *in* $> 3d(s_0)$ *than she promised in the previous configuration.*

At configuration $\langle s_0, [\alpha]_{>3d(s_0)}, \mathtt{V}\rangle$, player V *avoids losing only by choosing a sub-distribution d that maps s_0 to 0 and all other states to positive mass as before, and for the same reasons. Similarly, $d(s_1) < 1/3$ has to hold. So although a new function d with a new value of ϵ may be chosen, the next configuration is still of the same type $\langle s_0, [\alpha]_{>p'}, \mathtt{V}\rangle$ with $p' > 1/2$. Thus, either the play is finite and so lost for player* V *as described above; or the play is infinite and so lost for player* V *by the acceptance conditions A1 on infinite plays.*

We conclude that player R *wins that game. A winning strategy for her from the initial configuration only needs to be specified for move M9:*

- *player* R *will never choose a configuration of form $\langle s_0, q, \mathtt{V}\rangle$, should such an opportunity arise*
- *whenever player* V *chooses sub-distribution d with $d(s_2) > 0$, player* R *will choose s_2*
- *otherwise, it must be the case that both $d(s_1)$ and $d(s_2)$ are positive; if $d(s_1) = 1/3$, player* R *chooses s_1*
- *if $d(s_1) \neq 1/3$, player* R *chooses s_0*

5 Winning strategies

We show that when a player can win game $\mathsf{G}_M(s, \phi)$ she can use winning strategies that are of a very specific type. In addition to being memoryless in the classical sense, they choose very structured distributions when re-visiting a state in a configuration with a strong or weak until operator.

As before we use the notion of strategy informally. A strategy is *memoryless* if the choices of its player depend solely on the current configuration, not on the finite history of configurations that preceded the current one in a play. In our games, there can be configurations of type $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}\rangle$ for the same state s and the same path formula α (e.g., $\psi_1 \,\mathsf{U}\, \psi_2$) but with different bounds $\bowtie p$. We show that it is enough to consider winning strategies which induce bounds that change monotonically, as defined below. Subsequently, for sub-distributions $d, d' \colon S \to [0, 1]$, we write

- $d \leq d'$ iff for all $s \in S$ we have $d'(s) \leq d(s)$
- $d' < d$ iff $d' \leq d$ and $d'(s) < d(s)$ for some $s \in S$

For a *locally monotone* strategy the choice of sub-distribution d at configuration $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}\rangle$ is monotone in $\bowtie p$, regardless of the history of a play.

Definition 4 (Locally Monotone Strategies) *A strategy σ for player* C *in game $\mathsf{G}_M(s, \phi)$ is* locally monotone *iff for any two configurations $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}\rangle$ and $\langle s, [\alpha]_{\bowtie p'}, \mathtt{C}\rangle$ that occur in plays consistent with σ (but not necessarily in the same play), where d and d' are the sub-distributions chosen according to σ at these two configurations (respectively), then $p \geq p'$ implies $d \geq d'$ and $p > p'$ implies $d > d'$.*

A *cyclically monotone* strategy is monotone on cyclic paths within single plays: its player can force a decrease or increase of the thresholds depending on the path formula and whether it is a V or R configuration.

Definition 5 (Cyclically Monotone Strategies) *A strategy σ for player* C *in game $\mathsf{G}_M(s, \phi)$ is* cyclically monotone *iff for any two configurations $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}'\rangle$ and $\langle s, [\alpha]_{\bowtie p'}, \mathtt{C}'\rangle$ that occur in this order on some play consistent with σ, then*

- *$\alpha = \psi_1 \,\mathsf{U}\, \psi_2$ and $\mathtt{C} = \mathtt{C}'$ imply $p' < p$,*
- *$\alpha = \psi_1 \,\mathsf{W}\, \psi_2$ and $\mathtt{C} = \mathtt{C}'$ imply $p' \leq p$,*
- *$\alpha = \psi_1 \,\mathsf{U}\, \psi_2$ and $!\mathtt{C} = \mathtt{C}'$ imply $p' \geq p$,*
- *$\alpha = \psi_1 \,\mathsf{W}\, \psi_2$ and $!\mathtt{C} = \mathtt{C}'$ imply $p' > p$.*

The existence of winning strategies implies the existence of winning strategies that are locally monotone and cyclically monotone.

Theorem 2 *For every game $\mathsf{G}_M(s, \phi)$, there exists a winning strategy for player* C *iff there exists a memoryless winning strategy for player* C *that is also locally monotone and cyclically monotone.*

Example 9 *The winning strategy for Refuter in Example 8 is locally monotone as Refuter never encounters a pair of configurations that need to be checked for local monotonicity. That strategy is also cyclically monotone: From configuration $\langle s_0, [q \,\mathsf{U}\, r]_{>p}, \mathtt{V}\rangle$ the only possible cycles leads to configurations $\langle s_0, [q \,\mathsf{U}\, r]_{>p'}, \mathtt{V}\rangle$. As explained already, Verifier is restricted to $d(s_2) = 0$ and $d(s_1) < 1/3$ or she loses in the next step. Let $p > 1/2$ and $\epsilon = 1/3 - d(s_1)$. It follows that $d(s_0) \geq 1/6 + (p - 1/2) + \epsilon$. Thus, in the next configuration $\langle s_0, [q \,\mathsf{U}\, r]_{>p'}, \mathtt{V}\rangle$ we have $p' \geq 1/2 + 3(p - 1/2) + 3\epsilon$. As $\epsilon > 0$ and $p - 1/2 > 0$ we have $p' > p$. Furthermore, if $p_1, p_2, \ldots$ is the sequence of bounds obtained in this manner, then $p_{i+2} - p_{i+1} > p_{i+1} - p_i$ for all $i \geq 1$.*

6 Discussion

Table 1 summarizes which PCTL sub-formulae that may cause infinite plays can always be coerced into finite plays with a winning strategy. For example, a strong until with strict bound is ensured to have a finite strategy and explore a finite portion of the game before going to subformulae, and similarly from a negated weak until with a non-strict bound. To determine whether a PCTL formula is won by means of such finite plays only, we can either convert it into "GreaterThan" normal form and check whether each such sub-formula has a negation polarity that corresponds to the desired player in that table, or we can convert it into negation normal form and interpret that table *as is* on the resulting sub-formulae. For example, formula

Table 1. Sub-formulae that result in finite plays ($\surd$) or don't (x), for which winning player; ticks in parentheses indicate finite plays after an initial ϵ-correction of bounds

	$\mathsf{X}_{>}$	$\mathsf{X}_{\geq}$	$\mathsf{W}_{>}$	$\mathsf{W}_{\geq}$	$\mathsf{U}_{>}$	$\mathsf{U}_{\geq}$
Verifier	$\surd$	x ($\surd$)	x	x	$\surd$	x ($\surd$)
Refuter	x	x	x($\surd$)	$\surd$	x	x

$[q \,\mathsf{U}\, r]_{>0.999} \wedge \neg[q \,\mathsf{W}\, r]_{\geq 0.9991}$ is such that player V can win be ensuring only finite plays, if she can win at all. Furthermore, if the Markov chain is infinite, the game explores only a finite portion of it. From a practical point of view, it may be possible to change the strictness of the bound by slightly changing the required probabilities in the formula. Thus, an ϵ-correction of the formula may change a formula that does not allow finite plays to a formula that does allow finite plays.

7 Conclusions

We captured the PCTL semantics over countably labeled Markov chains through Hintikka games with Büchi acceptance conditions. Games moves depend on the strictness or non-strictness of probability thresholds for path formulae. Winning strategies may be assumed to be memoryless and monotone in their choice of structural elements (here sub-distributions). PCTL formulae in "GreaterThan" normal form that contain until operators with a certain combination of threshold type and negation polarity – statically derived from Table 1 – have winning strategies that may be interpreted as a finite-state abstraction of the underlying model that witnesses the falsity (respectively, truth) of the formula under consideration.

Acknowledgments. This research was in part supported by the UK EPSRC projects *Efficient Specification Pattern Library for Model Validation* (EP/D50595X/1) and *Complete and Efficient Checks for Branching-Time Abstractions* (EP/E028985/1).

References

[1] IEEE standard for a high performance serial bus, August 1996. Std 1394-1995.

[2] H. Aljazzar and S. Leue. Counterexamples for model checking of markov decision processes. Technical Report soft-08-01, University of Konstanz, December 2007. abstract.

[3] A. Condon. The complexity of stochastic games. *Information and Computation*, 96:203–224, 1992.

[4] D. Dams and K. Namjoshi. The existence of finite abstractions for branching time model checking. In *Logic in Computer Science*, 2004.

[5] D. Dams and K. Namjoshi. Automata as abstractions. In *International conference on Verification, Model Checking, and Abstract Interpretation*, volume 3385 of *Lecture Notes in Computer Science*, pages 216–232. Springer-Verlag, 2005.

[6] J. Desharnais, A. Edalat, and P. Panangaden. Bisimulation for Labelled Markov Processes. *Information and Computation*, 179:163–193, 2002.

[7] J. Desharnais, R. Jagadeesan, V. Gupta, and P. Panangaden. Approximating labeled Markov processes. In *Proceedings of the 15th Annual IEEE Symposium on Logic in Computer Science*, pages 95–106. IEEE, 2000.

[8] J. Edmund M. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. The MIT Press, 1999.

[9] E. Emerson and C. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *Foundations of Computer Science*, pages 368–377. IEEE, 1991.

[10] D. Fischer, E. Grädel, and L. Kaiser. Model checking games for the quantitative μ-calculus. In *Dans Proceedings of the 25th Annual Symposium on the Theoretical Aspects of Computer Science (STACS 2008)*, pages 301–312. arXiv:0802.2871v1 [cs.LO], 2008.

[11] T. Han and J.-P. Katoen. Counterexamples in probabilistic model checking. In *Tools and Algorithms for the Construction and Analysis of Systems*, 2007.

[12] H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6:512–535, 1994.

[13] J. Hintikka. *Logic, Language-Games and Information: Kantian Themes in the Philosophy of Logic*. Clarendon Press, Oxford, 1973.

[14] J. G. Kemeny, J. L. Snell, and A. W. Knapp. *Denumerable Markov Chains*. Springer Verlag, 1976. Second Edition.

[15] D. Kozen. Semantics of probabilistic programs. *Journal of Computer and Systems Sciences*, 22:328–350, 1981.

[16] D. Kozen. A probabilistic PDL. In *Proceedings of the fifteenth annual ACM Symposium on Theory of Computing*, 1983.

[17] D. Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27:333–354, 1983.

[18] M. Kwiatkowska, G. Norman, and D. Parker. Game-based abstraction for Markov decision processes. In *Proceedings of the 3rd international conference on the Quantitative Evaluation of Systems*, pages 157–166. IEEE, 2006.

[19] K. G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94:1–28, 1991.

[20] D. A. Martin. Borel Determinacy. *Annals of Mathematics*, 102:363–371, 1975.

[21] C. Morgan and A. McIver. Results on the quantitative μ-calculus qmμ. *ACM Transactions on Computational Logic (TOCL)*, 8(1), 2007.

[22] A. D. Pierro, C. Hankin, and H. Wiklicky. Measuring the confinement of probabilistic systems. *Theoretical Computer Science*, 340(1):3–56, 2005.

[23] T. Wilke. Alternating tree automata, parity games, and modal μ-calculus. *Bull. Soc. Math. Belg.*, 8(2), May 2001.

A Proofs

Proof of Lemma 1.

Consider first the case that M is finitely branching. Recall that $\mathsf{Path}(s, q \,\mathsf{U}\, r)$ denotes the set of paths beginning in s that satisfy $q \,\mathsf{U}\, r$. Let $\mathsf{Path}_i(s, q \,\mathsf{U}\, r)$ be $\mathsf{Path}(s, (q \,\mathsf{U}^{\leq i} r) \wedge \bigwedge_{0 \leq j < i} \neg(q \,\mathsf{U}^{\leq j} r))$, i.e., paths in which q holds until location i where r holds and r does not hold in locations smaller than i. We set $\mathsf{Path}_0(s, q \,\mathsf{U}\, r)$ to be $\mathsf{Path}(s, q \,\mathsf{U}^{\leq 0} r)$, i.e. the set $\{\pi = s_0 \cdots \mid s = s_0,\ s_0 \in L(r)\}$.

- For the "if" part, assume that for all $n \in \mathbb{N}$ there is $k \geq 0$ such that $s \in [\![[q \,\mathsf{U}\, r]_{>p-1/n}]\!]_{M^s_k}$. Then, $s \in [\![[q \,\mathsf{U}\, r]_{>p-1/n}]\!]_M$ follows by the monotonicity of the denotational semantics for "GreaterThan" thresholds. Thus, $s \in \bigcap_{n \in \mathbb{N}} [\![[q \,\mathsf{U}\, r]_{>p-1/n}]\!]_M = [\![[q \,\mathsf{U}\, r]_{\geq p}]\!]_M$.
- For the "only if" part, let $s \in [\![[q \,\mathsf{U}\, r]_{\geq p}]\!]_M$ and $n \in \mathbb{N}$. It suffices to find some $k \geq 0$ with $s \in [\![[q \,\mathsf{U}\, r]_{>p-1/n}]\!]_{M^s_k}$. As $\mathsf{Path}_i(s, q \,\mathsf{U}\, r)$ is of form $\mathsf{Path}(s, \alpha)$ for a path formula α, that set of paths is measurable. For all $i \neq j$ note that sets $\mathsf{Path}_i(s, q \,\mathsf{U}\, r)$ and $\mathsf{Path}_j(s, q \,\mathsf{U}\, r)$ are disjoint. Since

$$\mathsf{Path}(s, q \,\mathsf{U}\, r) = \bigcup_{i \geq 0} \mathsf{Path}_i(s, q \,\mathsf{U}\, r)$$

and as the latter is a disjoint union, we know that

$$\mathsf{Prob}_M(s, \mathsf{Path}(s, q \,\mathsf{U}\, r)) = \Sigma_{i \geq 0} \mathsf{Prob}_M(s, \mathsf{Path}_i(s, q \,\mathsf{U}\, r))$$

By definition of convergence for that infinite sum, for every $n \in \mathbb{N}$ there exists $k \geq 0$ such that

$$\Sigma^k_{i=0} \mathsf{Prob}_M(s, \mathsf{Path}_i(s, q \,\mathsf{U}\, r)) \geq \mathsf{Prob}_M(s, \mathsf{Path}(s, q \,\mathsf{U}\, r)) - 1/n$$

As $\Sigma^k_{i=1} \mathsf{Prob}_M(s, \mathsf{Path}_i(s, q \,\mathsf{U}\, r))$ equals $\mathsf{Prob}_{M^s_k}(s, q \,\mathsf{U}\, r)$ we obtain $s \in [\![[q \,\mathsf{U}\, r]_{>p-1/n}]\!]_{M^s_k}$ and we are done.

As M is finitely branching, there exists l such that l is an upper bound on the branching degree for all states in M^s_k. It follows that $\mathsf{Prob}_{M^s_k}(s, q \,\mathsf{U}\, r) = \mathsf{Prob}_{M^s_{k,l}}(s, q \,\mathsf{U}\, r)$.

In the case that M has infinite branching the proof is similar. We have to be more careful in noticing that every path set $\mathsf{Path}_i(s, q \,\mathsf{U}\, r)$ is still measurable and have to be careful in the way in which we sum up the probability of the set $\mathsf{Path}(s, q \,\mathsf{U}\, r)$. But this works out since all infinite sums have absolute convergence, establishing that for some k we have $s \in [\![[q \,\mathsf{U}\, r]_{>p-1/n}]\!]_{M^s_k}$. The existence of $M^s_{k,l}$ as required follows from convergence of $\mathsf{Prob}_{M^s_{k,l}}(s, q \,\mathsf{U}\, r)$ to $\mathsf{Prob}_{M^s_k}(s, q \,\mathsf{U}\, r)$. □

Corollary 1 *For labeled Markov chain $M = (S, P, L)$, $q, r \in \mathbb{AP}$, and $p \in [0, 1]$: $s \notin [\![[q \,\mathsf{W}\, r]_{>p}]\!]_M$ iff for all $n \in \mathbb{N}$ there are $k, l \in \mathbb{N}$ with $s \notin [\![[q \,\mathsf{W}\, r]_{\geq p+1/n}]\!]_{M^s_{k,l}}$.*

Proof of Corollary 1.

This follows from Lemma 1 and the duality of weak and strong until. We have $s \notin [\![[q \,\mathsf{W}\, r]_{>p}]\!]_M$ iff $s \in [\![[\neg r \,\mathsf{U}\, (\neg q \wedge \neg r)]_{\geq 1-p}]\!]_M$, since (semantically) $\neg r \,\mathsf{U}\, (\neg q \wedge \neg r)$ is the negation of $q \,\mathsf{W}\, r$. By Lemma 1, $s \in [\![[\neg r \,\mathsf{U}\, (\neg q \wedge \neg r)]_{\geq 1-p}]\!]_M$ holds iff for every $n \in \mathbb{N}$ there is $k, l \in \mathbb{N}$ with $s \in [\![[\neg r \,\mathsf{U}\, (\neg q \wedge \neg r)]_{>1-p-1/n}]\!]_{M^s_{k,l}}$, i.e. $s \notin [\![[q \,\mathsf{W}\, r]_{\geq p+1/n}]\!]_{M^s_{k,l}}$. □

Proof of Theorem 1.

Given PCTL formula ϕ, we show these two items by structural induction on the PCTL formulae ψ in the closure of ϕ, simultaneously on all states of M. Since exactly one of $s \in [\![\psi]\!]_M$ and $s \notin [\![\psi]\!]_M$ holds, it suffices to show both items in Theorem 1 for such a ψ in their "only if" versions, which we do by splitting ψ into six cases:

Case #1. The cases when ψ equals $\mathtt{tt}$ or $\mathtt{ff}$ are trivial. For example, no state satisfies $\mathtt{ff}$ and all plays beginning in $\langle s, \mathtt{ff}, \mathtt{V}\rangle$ are won by player $\mathtt{R}$. *So we may implicitly assume in subsequent cases that $\bowtie p$ equals neither > 1 nor ≥ 0.*

Case #2. The cases when ψ equals q, $\neg\psi_1$, or $\psi_1 \wedge \psi_2$ are proved as in the case of Hintikka semantic games for propositional logic. We illustrate this for the case of conjunction and player $\mathtt{V}$:

- Let $s \in [\![\psi_1 \wedge \psi_2]\!]_M$. Then $s \in [\![\psi_i]\!]_M$ for all $i = 1, 2$. By induction, there is a winning strategy w_i for player $\mathtt{V}$ from configuration $\langle s, \psi_i, \mathtt{V}\rangle$ for all $i = 1, 2$. Consider the strategy w for player $\mathtt{V}$ which composes his strategies w_1 and w_2 as follows: at configuration $\langle s, \psi, \mathtt{V}\rangle$, player $\mathtt{R}$ has to choose as next configuration some $\langle s, \psi_i, \mathtt{V}\rangle$ with $i = 1, 2$. But then player $\mathtt{V}$ simply responds according to his winning strategy w_i. This describes a winning strategy w for player $\mathtt{V}$ from configuration $\langle s, \psi, \mathtt{V}\rangle$.
- Let $s \notin [\![\psi_1 \wedge \psi_2]\!]_M$. Then there is some $j \in \{1, 2\}$ such that $s \notin [\![\psi_j]\!]_M$. By induction, there is a winning strategy w_j for player $\mathtt{R}$ from configuration $\langle s, \psi_j, \mathtt{V}\rangle$. Consider the strategy w for player $\mathtt{R}$ which composes his strategy w_j with an initial choice as follows: at configuration $\langle s, \psi, \mathtt{V}\rangle$, player $\mathtt{R}$ simply chooses $\langle s, \psi_j, \mathtt{V}\rangle$ as next configuration and then plays according to his winning strategy w_j. This describes a winning strategy w for player $\mathtt{R}$ from configuration $\langle s, \psi, \mathtt{V}\rangle$.

Case #3. The case when ψ equals $[\mathsf{X}\, \psi_1]_{\bowtie p}$, where $\bowtie \in \{\geq, >\}$:

- Let $s \in [\![[\mathsf{X}\,\psi_1]_{\bowtie p}]\!]_M$. Let Y be the set of states s' such that $P(s,s') > 0$ and $s' \in [\![\psi_1]\!]_M$. From the latter and induction we infer that player V has a winning strategy $w_{s'}$ for the configuration $\langle s', \psi_1, \mathtt{V}\rangle$, for all $s' \in Y$. We construct from all of these $w_{s'}$ a winning strategy w for player V from configuration $\langle s, [\mathsf{X}\,\psi_1]_{\bowtie p}, \mathtt{V}\rangle$ as follows: Since $s \in [\![[\mathsf{X}\,\psi_1]_{\bowtie p}]\!]_M$, we know that $P(s,Y) \bowtie p$ holds and that Y is non-empty as $\bowtie p$ isn't ≥ 0. So at configuration $\langle s, [\mathsf{X}\,\psi_1]_{\bowtie p}, \mathtt{V}\rangle$ player V chooses this set Y. Now no matter what next configuration $\langle s', \psi_1, \mathtt{V}\rangle$ player R chooses, we have $s' \in Y$ and so player V will play according to his winning strategy $w_{s'}$. In particular, w is a winning strategy for player V from configuration $\langle s, [\mathsf{X}\,\psi_1]_{\bowtie p}, \mathtt{V}\rangle$
- Let $s \notin [\![[\mathsf{X}\,\psi_1]_{\bowtie p}]\!]_M$. Player V must choose a set Y such that $P(s,Y) \bowtie p$. In making this choice, player V would – by induction – lose from configuration $\langle s, [\mathsf{X}\,\psi_1]_{\bowtie p}, \mathtt{V}\rangle$ if Y contained some s' with $s' \notin [\![\psi_1]\!]_M$ (for then player R could respond with configuration $\langle s', \psi_1, \mathtt{V}\rangle$ and win the resulting game). Dually, player V can only increase her chances of winning from configuration $\langle s, [\mathsf{X}\,\psi_1]_{\bowtie p}, \mathtt{V}\rangle$ if she adds to Y all states s' with $s' \in [\![\psi_1]\!]_M$ and $P(s,s') > 0$. Finally, player V has no incentive to add an $s' \in [\![\psi_1]\!]_M$ to Y if $P(s,s') = 0$: this does not contribute to $P(s,Y) \bowtie p$ and only exposes player V to a threat of player R to move to s'. To summarize, if player V has a winning strategy from that configuration, then she also has a winning strategy from that same configuration where she chooses Y as in the previous item. But then $s \notin [\![[\mathsf{X}\,\psi_1]_{\bowtie p}]\!]_M$ means that $P(s,Y) \bowtie p$ is false. So player V can only choose a set Y for which player R can respond with a winning strategy.

Case #4. The cases when ϕ equals $[\alpha]_{\geq p}$ where α is $\psi_1 \,\mathsf{U}\, \psi_2$:

- Let $s \in [\![\phi]\!]_M$. Then $\mathsf{Prob}_M(s,\alpha) \geq p$ and so for each $n \in \mathbb{N}$ with $p - 1/n \geq 0$ we have $\mathsf{Prob}_M(s,\alpha) > p - 1/n$ and $s \in [\![[\alpha]_{>p-1/n}]\!]_M$. By induction, player V has a winning strategy w_n from configuration $\langle s, [\alpha]_{>p-1/n}, \mathtt{V}\rangle$, for each such $n \in \mathbb{N}$. Player V can synthesize from these countably many strategies a winning strategy w for her from configuration $\langle s, \phi, \mathtt{V}\rangle$ as follows: if player R chooses any such $n \in \mathbb{N}$, then the next configuration is $\langle s, [\alpha]_{>p-1/n}, \mathtt{V}\rangle$ and player V plays according to w_n.
- Let $s \notin [\![\phi]\!]_M$. Then $\mathsf{Prob}_M(s,\alpha) < p$. Thus, there is some $n_0 \in \mathbb{N}$ with $\mathsf{Prob}_M(s,\alpha) \leq p - 1/n_0 < 1$. But then $\mathsf{Prob}_M(s,\alpha) \not> p - 1/n_0$ implies $s \notin [\![[\alpha]_{>p-1/n_0}]\!]_M$. By induction, player R has a winning strategy w_{n_0} from configuration $\langle s, [\alpha]_{>p-1/n_0}, \mathtt{V}\rangle$. So player R gets a winning strategy w from configuration $\langle s, \phi, \mathtt{V}\rangle$ by first choosing that n_0 and then playing according to w_{n_0}.

Case #5. The cases when ϕ equals $[\alpha]_{>p}$ where α is $\psi_1 \,\mathsf{W}\, \psi_2$:

- Let $s \in [\![\phi]\!]_M$. Then $\mathsf{Prob}_M(s,\alpha) > p$. Thus, there is some $n_0 \in \mathbb{N}$ with $p + 1/n_0 \leq 1$ and $\mathsf{Prob}_M(s,\alpha) \geq p + 1/n_0$. But then $\mathsf{Prob}_M(s,\alpha) \geq p + 1/n_0$ implies $s \in [\![[\alpha]_{\geq p+1/n_0}]\!]_M$. By induction, player V has a winning strategy w_{n_0} from configuration $\langle s, [\alpha]_{\geq p+1/n_0}, \mathtt{V}\rangle$. So player V gets a winning strategy w from configuration $\langle s, \phi, \mathtt{V}\rangle$ by first choosing that n_0 and then playing according to w_{n_0}.
- Let $s \notin [\![\phi]\!]_M$. Then $\mathsf{Prob}_M(s,\alpha) \leq p$. Thus, for every $n \in \mathbb{N}$ with $p + 1/n \leq 1$ we have $\mathsf{Prob}_M(s,\alpha) < p + 1/n$. By induction, player R has a winning strategy w_n from configuration $\langle s, [\alpha]_{>p+1/n}, \mathtt{V}\rangle$, for each $n \in \mathbb{N}$ with $p + 1/n \leq 1$. Player R can synthesize from these countable strategies a winning strategy for her from configuration $\langle s, \phi, \mathtt{V}\rangle$ as follows: if player V chooses such an $n \in \mathbb{N}$, then the next configuration is $\langle s, [\alpha]_{>p+1/n}, \mathtt{V}\rangle$ and player R plays according to w_n.

Case #6. The cases when ϕ equals $[\alpha]_{\bowtie p}$ where either

(a) α is $\psi_1 \,\mathsf{U}\, \psi_2$ and $\bowtie$ is $>$
(b) α is $\psi_1 \,\mathsf{W}\, \psi_2$ and $\bowtie$ is $\geq$ or
(c) α is $\psi_1 \,\mathsf{U}^{\leq k} \psi_2$ or $\psi_1 \,\mathsf{W}^{\leq k} \psi_2$ with $k \in \mathbb{N}$ and $\bowtie$ is either $>$ or $\geq$:

- Let $s \in [\![\phi]\!]_M$.

  The formula α is logically equivalent to $\psi_2 \vee (\psi_1 \wedge \mathsf{X}\,\alpha)$ and in case that α is bounded the bound decreases by 1. It follows that it is either the case that $s \in [\![\psi_2]\!]_M$ or $s \in [\![\psi_1 \wedge [\mathsf{X}\,\alpha]_{\bowtie p}]\!]_M$. In the first case, player V chooses to move to configuration $\langle s, \psi_2, \mathtt{V}\rangle$ and by induction she has a winning strategy from this configuration. In the second case, by induction there is a winning strategy for player V from configuration $\langle s, \psi_1, \mathtt{V}\rangle$, so if player R chooses to go to this configuration, player V wins. If player R does not move to ψ_1, then M9 demands that player V chooses a subdistribution $d : S \to [0,1]$ satisfying (4)-(5). By assumption $s \in [\![[\mathsf{X}\,\alpha]_{\bowtie p}]\!]_M$. Let T be the set of states t such that $\mathsf{Prob}_M(t,\alpha) > 0$ and $P(s,t) > 0$. We choose d such that $d(s') = 0$ for all $s' \in S \setminus T$.

  So it suffices to specify d on set T. For that, let $p' = \Sigma_{t \in T} P(s,t) \cdot \mathsf{Prob}_M(t,\alpha)$.

  – Consider the case that $\bowtie$ is $>$. By assumption $p' > p$. In the case that $p = 0$, we choose some state $t \in T$ such that $\mathsf{Prob}_M(t,\alpha) > 0$, we set $d(t) = \mathsf{Prob}_M(t,\alpha) \cdot P(s,t)$, and $d(t') = 0$ for

all $t' \neq t$. In the case that $p > 0$, let δ be $p' - p$. We are going to distribute this gap δ between all the states in T according to the distribution $P(s,\cdot)$. That is, for all $t \in T$

$$d(t) = max(0, (\mathsf{Prob}_M(t,\alpha) - \delta) \cdot P(s,t))$$

In case that $\mathsf{Prob}_M(t,\alpha) \leq \delta$ we thus have $d(t) = 0$ (and so effectively remove t from set T above). As $p' = \Sigma_{t \in S} \mathsf{Prob}_M(t,\alpha) P(s,t)$ and $p > 0$ there must be at least one state t such that $\mathsf{Prob}_M(t,\alpha) \geq p'$ and hence $\mathsf{Prob}_M(t,\alpha) - \delta > 0$, implying $d(t) > 0$. It follows that $\Sigma_{t \in T} d(t) \geq p' - \delta \geq p$.

– Consider the case that $\bowtie$ is $\geq$. By assumption $p' \geq p$. Let δ be $p' - p$. For all $t \in T$, let

$$d(t) = max(0, \mathsf{Prob}_M(t,\alpha) - \delta \cdot P(s,t))$$

Again, if $\mathsf{Prob}_M(t,\alpha) \leq \delta$ we set $d(t) = 0$. This completes the specification of sub-distribution d chosen by player V.

Now regardless of the choice of player R, the next configuration is $\langle t, [\alpha]_{\bowtie p'}, \mathtt{V}\rangle$ such that $t \in [\![[\alpha]_{\bowtie p'}]\!]_M$. So player V maintains the truth value of the configuration. Notice that also the distance from the promised bound p' and the real probability is being maintained.

Case (c): For bounded operators, as the bound decreases, in a finite number of steps the play moves to configurations of the form $\langle s', \psi_i, \mathtt{V}\rangle$ for $i \in \{1,2\}$, where induction applies directly, and in the desired manner.

Case (b): For Weak Until $\psi_1 \,\mathsf{W}\, \psi_2$, all infinite plays have a suffix of configurations of form $\langle s', [\psi_1 \,\mathsf{W}\, \psi_2]_{\geq p}, \mathtt{V}\rangle$ and are thus winning for player V. Finite plays again reach configurations of the form $\langle s', \psi_i, \mathtt{V}\rangle$ for $i \in \{1,2\}$, where induction applies directly.

Case (a): For (strong) Until, we appeal to Lemma 1. We treat subformulae ψ_1 and ψ_2 as propositions (respectively, the q and r in that lemma) and annotate states of M by ψ_1 and ψ_2. Let $p' = \mathsf{Prob}_M(s, \psi_1 \,\mathsf{U}\, \psi_2)$. By assumption $p' > p$. In particular, $s \in [\![[\psi_1 \,\mathsf{U}\, \psi_2]_{\geq p'}]\!]_M$. Let $n \in \mathbb{N}$ be such that $p' > p' - 1/n > p$. By Lemma 1 (applied to p' instead of p), there are $k, l \geq 0$ with $s \in [\![[\psi_1 \,\mathsf{U}\, \psi_2]_{>p'-1/n}]\!]_{M^s_{k,l}}$ and so the probability of $\psi_1 \,\mathsf{U}\, \psi_2$ in $M^s_{k,l}$ at s is greater than p. Player V's strategy is to consider this system $M^s_{k,l}$. She chooses sub-distributions $d\colon S \to [0,1]$ according to the probabilities $\mathsf{Prob}_{M^s_{k,l}}(t,\alpha)$ (instead of $\mathsf{Prob}_M(t,\alpha)$ but as explained above). By definition of $M^s_{k,l}$ there can be only finite sequences of configurations of the form $\langle s', [\alpha]_{>p}, \mathtt{V}\rangle$, and so player V wins (cf. Example 7).

- Let $s \notin [\![\phi]\!]_M$.

  It follows that $\mathsf{Prob}_M(s,\alpha) \leq p$ in case that $\bowtie$ is $>$; and $\mathsf{Prob}_M(s,\alpha) < p$ in case that $\bowtie$ is $\geq$. As above, α is logically equivalent to $\psi_2 \vee (\psi_1 \wedge \mathsf{X}\,\alpha)$ and in case that α is bounded the bound decreases by 1. It follows that $s \notin [\![\psi_2]\!]_M$ and hence there is a winning strategy for player R from configuration $\langle s, \psi_2, \mathtt{V}\rangle$. Also, it is either the case that $s \notin [\![\psi_1]\!]_M$ or $s \notin [\![[\mathsf{X}\,\alpha]_{\bowtie p}]\!]_M$. In the first case player R has a winning strategy from configuration $\langle s, \psi_1, \mathtt{V}\rangle$ and chooses this configuration. In the second case, player V chooses a sub-distribution $d\colon S \to [0,1]$ such that (4)-(5) hold.

  We claim that there is some $s' \in S$ with $d(s') > 0$ and $\mathsf{Prob}_M(s',\alpha) \not\bowtie d(s') P(s,s')^{-1}$. Proof by contradiction: otherwise, $\mathsf{Prob}_M(s',\alpha) \bowtie d(s')$ for all s' with $d(s') > 0$ implies that

$$\sum_{s' | d(s') > 0} \mathsf{Prob}_M(s',\alpha) \bowtie \sum_{s' \in S} d(s') \geq p$$

  by (4). But this renders

$$\sum_{s' | d(s') > 0} \mathsf{Prob}_M(s',\alpha) \bowtie p$$

  which directly contradicts $s \notin [\![[\mathsf{X}\,\alpha]_{\bowtie p}]\!]_M$.

  Thus, player R can choose such an s' and maintain the play in configurations of the form $\langle s', [\alpha]_{\bowtie p'}, \mathtt{V}\rangle$ such that $s' \notin [\![[\alpha]_{\bowtie p'}]\!]_M$. Notice that player R can choose a successor s' such that

$$p' - \mathsf{Prob}_M(s',\alpha) \geq p - \mathsf{Prob}_M(s,\alpha)$$

  i.e., the gap between the promise and the actual probability does not decrease.

  We now study the consequences of this capability of player R for the different forms of path formula α in this case:

  Case (c): For bounded operators, as the bound decreases, in a finite number of steps the play moves to configurations of the form $\langle s', \psi_i, \mathtt{V}\rangle$ for $i \in \{1,2\}$ and so player R wins by induction.

  Case (b): For (strong) Until formulae, infinite plays of configurations of the form $\langle s', [\psi_1 \,\mathsf{U}\, \psi_2]_{\bowtie p}, \mathtt{V}\rangle$ are winning for player R by the winning conditions for infinite plays. Any finite play reduces to configurations of the form $\langle s', \psi_i, \mathtt{V}\rangle$ for $i \in \{1,2\}$, where induction applies directly, and in the desired manner.

  Case (a): For Weak Until formulae, we appeal to Corollary 1. As before, we treat ψ_1 and ψ_2 as propositions and annotate states of M by them. Let $p' = \mathsf{Prob}_M(s, \psi_1 \,\mathsf{W}\, \psi_2)$. By assumption $p' \leq p$. In particular, $s \notin [\![[\psi_1 \,\mathsf{W}\, \psi_2]_{>p'}]\!]_M$. Let $n \in \mathbb{N}$ be such

that $p' < p + 1/n < p$. By Corollary 1 there are $k, l \geq 0$ with $s \notin [\![[\psi_1 \mathsf{W} \psi_2]_{\geq p'+1/n}]\!]_{M^s_{k,l}}$ and so the probability of $\psi_1 \mathsf{W} \psi_2$ in $M^s_{k,l}$ at s is less than p. Player R's strategy is to consider this system $M^s_{k,l}$. Let $d\colon S \to [0,1]$ be the sub-distribution chosen by player V. As $s \notin [\![[\psi_1 \mathsf{W} \psi_2]_{\geq p}]\!]_{M^s_{k,l}}$, there is some $s' \in S$ such that $s' \notin [\![[\psi_1 \mathsf{W} \psi_2]_{\geq d(s')P(s,t)^{-1}}]\!]_{M^s_{k,l}}$. So player R chooses this s'. By definition of $M^s_{k,l}$ there can be only finite sequence of configuration of the form $\langle s', [\alpha]_{\geq p}, \mathtt{V}\rangle$, and so player R wins. This is dual to the strategy depicted for V in Example 7.

□

Proof of Theorem 2.

Assuming that there exists some winning strategy for player C in game $\mathsf{G}_M(s,\phi)$, it suffices to show that a slight modification of the winning strategy synthesized in the proof of Theorem 1 is memoryless, locally monotone, and cyclically monotone. That slightly modified strategy will clearly be memoryless by construction. We now describe this modified winning strategy and first prove its local monotonicity, by induction as in the proof of Theorem 1. Then we prove that it is cyclically monotone.

Modified winning strategy and its local monotonicity. The only configurations where player C needs to make choices are $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}'\rangle$, $\langle s, \psi_1 \vee \psi_2, \mathtt{C}\rangle$, and $\langle s, \psi_1 \wedge \psi_2, !\mathtt{C}\rangle$.

With the latter two, we restrict C's strategy to choose ψ_1 whenever possible and only when impossible choose ψ_2. This is similar to what one can do in Hintikka games for first-order logic. We show that the way configurations of the form $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}'\rangle$ are handled induces a memoryless and monotone strategy.

If $\alpha = X\psi$, then the strategy defined in the proof of Theorem 1 chooses the set of successors according to the state s, and is clearly memoryless.

If $!\mathtt{C} = \mathtt{C}'$ and either $\alpha = \psi_1 \mathsf{U} \psi_2$ and $\bowtie \; = \; \geq$ or $\alpha = \psi_1 \mathsf{W} \psi_2$ and $\bowtie \; = \; >$, then player C has to choose a value n. By choosing the minimal possible n she ensures that the strategy is memoryless.

Consider two configurations $\langle s, [\alpha]_{\bowtie p_1}, \mathtt{C}\rangle$ and $\langle s, [\alpha]_{\bowtie p_2}, \mathtt{C}\rangle$. Whenever the play moves to configurations of the form $\langle s', \psi_i, \mathtt{V}\rangle$ for $i \in \{1,2\}$, the strategy is memoryless, locally monotone, and cyclically monotone by induction. We start with proving local monotonicity for moves that may choose sub-distributions.

1. For configurations where $\alpha = \psi_1 \mathsf{W} \psi_2$, $\alpha = \psi_1 \mathsf{W}^{\leq k} \psi_2$, or $\alpha = \psi_1 \mathsf{U}^{\leq k} \psi_2$, and $\mathtt{C} = \mathtt{C}'$ we claim that the strategy composed in the proof of Theorem 1 is locally monotone by induction. Intuitively, this can be seen by the strategy using the gap δ between the probability of the formula and the required threshold. The strategy partitions this gap between all successors, so if the same state is visited with different thresholds the partition of the gap implies that the distribution decreases.

Let $p' = \mathsf{Prob}_M(s,\alpha)$ and $\delta_i = p' - p_i$ for $i \in \{1,2\}$. According to the proof of Theorem 1 in configuration $\langle s, [\alpha]_{\bowtie p_i}, \mathtt{C}\rangle$ player C chooses the distribution

$$d_i(t) = \max(0, (\mathsf{Prob}_M(t,\alpha) - \delta_i) \cdot P(s,t))$$

It follows that if $p_1 \geq p_2$ then for every $t \in S$ we have $d_1(t) \geq d_2(t)$. If follows that if $p_1 = p_2$ then $d_1 = d_2$. Consider the case that $p_1 > p_2$ Then $p_1 > 0$ and for some t we have $d_1(t) > 0$ and $d_1(t) = \mathsf{Prob}_M(t,\alpha) - \delta_1$. As $\delta_1 < \delta_2$ and $d_2(t) = \mathsf{Prob}_M(t,\alpha) - \delta_2$ it follows that $d_1(t) > d_2(t)$.

2. For the case where $\alpha = \psi_1 \mathsf{U} \psi_2$ and $\mathtt{C} = \mathtt{C}'$, the strategy as defined in the proof of Theorem 1 is not locally monotone. We modify it as follows: For every configuration $\langle s, [\psi_1 \mathsf{U} \psi_2]_{>p}, \mathtt{C}\rangle$ the sub-distribution d is chosen according to the minimal k such that some fraction of $\mathsf{Prob}_{M^s_k}(s,\alpha)$ is greater than p. The exact definition of this fraction is given below. Furthermore, we use the gap between $\mathsf{Prob}_{M^s_k}(s,\alpha)$ and $\mathsf{Prob}_{M^s_{k-1}}(s,\alpha)$ to ensure local (and later cyclic) monotonicity. The definition of the sub-distribution d and the proof itself are quite technical.

Consider the configuration $\langle s, [\alpha]_{>p}, \mathtt{C}\rangle$. We assume, without loss of generality, that $s \notin [\![\psi_2]\!]_M$. We measure the exact probability to satisfy α within i steps. For every $t \in S$ let

$$\begin{aligned} n^t_0 &= \mathsf{Prob}_{M^t_0}(t,\alpha) \\ n^t_i &= \mathsf{Prob}_{M^t_i}(t,\alpha) - \mathsf{Prob}_{M^t_{i-1}}(t,\alpha) \end{aligned}$$

Consider the following increasing sequence:

$$\begin{aligned} N^t_0 &= \frac{n^t_0}{2} \\ N^t_i &= N^t_{i-1} + \Sigma^i_{j=0} \frac{1}{2^{i+1-j}} n^t_j \qquad (i > 0) \end{aligned}$$

That is, $N^t_1 = \frac{3}{4}n^t_0 + \frac{1}{2}n^t_1$, $N^t_2 = \frac{7}{8}n^t_0 + \frac{3}{4}n^t_1 + \frac{1}{2}n^t_2$, $N^t_3 = \frac{15}{16}n^t_0 + \frac{7}{8}n^t_1 + \frac{3}{4}n^t_2 + \frac{1}{2}n^t_1$, and so on. Notice that

$$\lim_{i\to\infty} N^t_i = \mathsf{Prob}_{M^t_k}(t,\alpha)$$

Let i_0 be the minimal such that

$$\Sigma_{t\in S} N^t_{i_0} P(s,t) > p$$

By abuse of notation for $i \geq 0$, we denote

$$N^s_{i+1} = \Sigma_{t\in S} N^t_i P(s,t)$$

That is, N_i^s is the sum of the different N_{i-1}^t normalized by their probabilities to get from s to t. To simplify notations, for $i < 0$ and for all t we set

$$N_i^t = N_{i+1}^s = 0$$

The value $N_{i_0}^t P(s,t)$ is going to be the basis for defining $d(t)$. Notice that it must be the case that $N_{i_0}^s \leq p$ and that $N_{i_0}^t - N_{i_0-1}^t > 0$. In order to maintain local monotonicity we distribute the gap between the required threshold p and $N_{i_0}^s$ between all the states t where $N_{i_0+1}^t > 0$. We have to be extremely careful with the states s for which $N_{i_0}^s = p$. For these states, we take a constant fraction of $N_{i_0}^t - N_{i_0-1}^t$ and distribute it among the successors t. We then have to scale the distribution d for all states s for which this constant fraction surpasses the required bound.
We set $d(t)$ as follows:

$$d(t) = \left(N_{i_0-1}^t + \left(\tfrac{1}{4} + \tfrac{3}{4}\frac{p - N_{i_0}^s}{N_{i_0+1}^s - N_{i_0}^s}\right)\left(N_{i_0}^t - N_{i_0-1}^t\right)\right) P(s,t)$$

It is simple to see that

$$\sum_{t \in S} d(t) > p$$

Indeed, $\sum_{t \in S} d(t)$ is the sum of the following three expressions:

$$\sum_{t \in S} N_{i_0-1}^t P(s,t) = N_{i_0}^s$$
$$\sum_{t \in S} \frac{N_{i_0}^t - N_{i_0-1}^t}{4} P(s,t) = \frac{N_{i_0+1}^s - N_{i_0}^s}{4}$$
$$\sum_{t \in S} \frac{3}{4}\frac{p - N_{i_0}^s}{N_{i_0+1}^s - N_{i_0}^s}(N_{i_0}^t - N_{i_0-1}^t)P(s,t) = \frac{3}{4}(p - N_{i_0}^s)$$

As $N_{i_0+1}^s > p$ the result follows.

Furthermore, when going to some successor t of s the choice of i_0 for s implies that for the choice of the sub-distribution d for t some value $i_0' < i_0$ is going to be used. Thus, the sequence of configurations of the form $\langle t', [\alpha]_{>p'}, \mathtt{C}\rangle$ is finite and player $\mathtt{C}$ is winning.

We show that this definition of the sub-distribution d implies local monotonicity. Consider two configurations $\langle s, [\alpha]_{>p_1}, \mathtt{C}\rangle$ and $\langle s, [\alpha]_{>p_2}, \mathtt{C}\rangle$. Let d_1 and d_2 be the sub-distributions chosen by σ in these configurations and let i_0^1 and i_0^2 be the values used to define d_1 and d_2, respectively. By definition $d_j(t)$ is in the open interval

$$(N_{i_0^j-1}^t P(s,t), N_{i_0^j}^t P(s,t))$$

for $j \in \{1,2\}$. By definition if $p_1 = p_2$ then $i_0^1 = i_0^2$ and it follows that $d_1 = d_2$. Similarly, if $p_1 > p_2$ then $i_0^1 \geq i_0^2$. If $i_0^1 > i_0^2$ the strictness of $d_1 > d_2$ follows from the strictness of the sequence N_i^t. If $i_0^1 = i_0^2$ then $d_1 > d_2$ as $p_1 > p_2$.

Cyclic monotonicity of modified winning strategy. We turn now to consider cyclic monotonicity. Consider the configurations $\langle s, [\alpha]_{\bowtie p_1}, \mathtt{C}'\rangle$ and $\langle s, [\alpha]_{\bowtie p_2}, \mathtt{C}'\rangle$ that appear in a play consistent with σ according to this order.

- Consider the case where $\alpha = \psi_1 \mathsf{W} \psi_2$, $\alpha = \psi_1 \mathsf{W}^{\leq k} \psi_2$, or $\alpha = \psi_1 \mathsf{U}^{\leq k} \psi_2$ and $\mathtt{C} = \mathtt{C}'$. The strategy defined in the proof of Theorem 1 is also cyclically monotone. Indeed, from configuration $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}\rangle$ where
$$\mathsf{Prob}_M(s,\alpha) - p = \delta$$
we pass to configuration $\langle t, [\alpha]_{\bowtie p'}, \mathtt{C}\rangle$ and we know that
$$\mathsf{Prob}_M(t,\alpha) - p' = \delta$$
Hence, if configurations $\langle s, [\alpha]_{\bowtie p_1}, \mathtt{C}\rangle$ and $\langle s, [\alpha]_{\bowtie p_2}, \mathtt{C}\rangle$ appear in the same play we have $p_1 \geq p_2$.
- Consider the case where $\alpha = \psi_1 \mathsf{U} \psi_2$ and $\mathtt{C} = \mathtt{C}'$ and the strategy defined above. Let i_0^1 be the bound used for choosing the sub-distribution d in configuration $\langle s, [\alpha]_{>p_1}, \mathtt{C}\rangle$. By construction values smaller than i_0^1 are going to be used to define the sub-distributions in successor configurations. It follows that if configuration $\langle s, [\alpha]_{>p_2}, \mathtt{C}\rangle$ is visited, a value $i_0^2 < i_0^1$ is going to be used to define its sub-distribution. From the strictness of the sequence N_i^t (and N_i^s) and as $N_{i_0^j}^s \leq p_j < N_{i_0^j+1}^s$ it follows that $p_2 < p_1$.
- Consider the case where $\alpha = \psi_1 \mathsf{U} \psi_2$, $\alpha = \psi_1 \mathsf{U}^{\leq k}$, or $\alpha = \psi_1 \mathsf{W}^{\leq k} \psi_2$ and $!\mathtt{C} = \mathtt{C}'$. Let $p' = \mathsf{Prob}_M(s', \alpha)$ and $\delta_i = p_i - p'$ for $i \in \{1,2\}$. Let d be the distribution suggested by player $!\mathtt{C}$ in configuration $\langle s, [\alpha]_{\bowtie p_1}, !\mathtt{C}\rangle$. By definition of d we have $\Sigma_{t \in S} d(t) \geq p_1$. By assumption $\langle s, [\alpha]_{\bowtie p_2}, !\mathtt{C}\rangle$ is reachable from $\langle s, [\alpha]_{\bowtie p_1}, !\mathtt{C}\rangle$, so both players do not choose to go to configurations of the form $\langle t, \psi_i, !\mathtt{C}\rangle$ for $i \in \{1,2\}$. If follows that
$$\mathsf{Prob}_M(s,\alpha) = \Sigma_{t \in S} P(s,t) \mathsf{Prob}_M(t,\alpha)$$
We know that
$$\Sigma_{t \in S} d(t) \geq p' + \delta_1$$
Then, there must exist some $t \in S$ such that
$$d(t) \cdot P(s,t)^{-1} \geq \mathsf{Prob}_M(t,\alpha) + \delta_1$$
It follows that if player $\mathtt{C}$ chooses this state t the gap between the actual probability and the threshold does not decrease. Thus $p_1 \leq p_2$.
- Consider the case where $\alpha = \psi_1 \mathsf{W} \psi_2$ and $!\mathtt{C} = \mathtt{C}'$. Then the proof is similar to the previous item. By assumption $\mathtt{C}$ wins from $\langle s, [\alpha]_{\geq p_1}, !\mathtt{C}\rangle$ and hence $s \notin$

$[\![[\alpha]_{\geq p_1}]\!]_M$. Let $p' = \mathsf{Prob}_M(s, \alpha)$. As player $\mathtt{C}$ wins from $\langle s, [\alpha]_{\geq p_1}, !\mathtt{C}\rangle$ we conclude that $p' < p_1$. In particular, $s \notin [\![[\psi_1 \,\mathsf{W}\, \psi_2]_{>p'}]\!]_M$. Let $n \in \mathbb{N}$ be such that $p' < p + 1/n < p$. By Corollary 1 there are $k, l \geq 0$ with $s \notin [\![[\psi_1 \,\mathsf{W}\, \psi_2]_{\geq p'+1/n}]\!]_{M^s_{k,l}}$ and so the probability of $\psi_1 \,\mathsf{W}\, \psi_2$ in $M^s_{k,l}$ at s is less than p_1. Player $\mathtt{C}$ is going to use system $M^s_{k,l}$ to guide her decisions. As usual

$$\mathsf{Prob}_{M^s_{k,l}}(s, \alpha) = \Sigma_{t \in S_{k,l}} P(s, t) \mathsf{Prob}_{M^s_{k,l}}(t, \alpha)$$

Let

$$p'' = \mathsf{Prob}_{M^s_{k,l}}(s, \alpha)$$

As mentioned $p'' < p_1$. Let $\delta_1 = p_1 - p''$ and let d be the distribution suggested by player $!\mathtt{C}$ in configuration $\langle s, [\alpha]_{\geq p_1}, !\mathtt{C}\rangle$. By definition of d we have

$$\Sigma_{t \in S} d(t) \geq p_1 = \delta_1 + p''$$

Then, there must exist some $t \in S$ such that

$$d(t) \cdot P(s, t)^{-1} \geq \mathsf{Prob}_{M^s_{k,l}}(t, \alpha) + \delta_1$$

It follows that if player $\mathtt{C}$ chooses this state t the gap between the actual probability in $M^s_{k,l}$ and the threshold does not decrease. We show below in Lemma 2 that when visiting the same state again in $M^s_{k,l}$ the probability of α increases. Hence, $p_2 > p_1$.

Lemma 2 *Let M be a labeled Markov chain, q and r in $\mathbb{AP}$, α the path formula $q \,\mathsf{W}\, r$, and $M^s_{k,l}$ given for some state s of M and $k, l \in \mathbb{N}$. Let t and t' be different states in $M^s_{k,l}$ that both correspond to some state s' of M such that*

- *there is a path from t to t' in t in $M^s_{k,l}$, and*
- *q holds throughout the unique and finite path from the root of $M^s_{k,l}$ to t'.*

If we have $\mathsf{Prob}_{M^s_k}(t, \alpha) < 1$, then $\mathsf{Prob}_{M^s_k}(t', \alpha) > \mathsf{Prob}_{M^s_k}(t, \alpha)$ follows.

Proof. As $\mathsf{Prob}_{M^s_k}(t, q \,\mathsf{W}\, r) < 1$ it follows that there is some "leaf" t'' in $M^s_{k,l}$ that is reachable from t in $M^s_{k,l}$ such that the unique finite path from t to t'' in $M^s_{k,l}$ does not satisfy $q \,\mathsf{W}\, r$. As $M^s_{k,l}$ is an unwinding of M, it follows that the subtree reachable from t' in $M^s_{k,l}$ is contained in the subtree reachable from t in $M^s_{k,l}$. Clearly, $\mathsf{Prob}_{M^s_{k,l}}(t', \alpha) \geq \mathsf{Prob}_{M^s_{k,l}}(t, \alpha)$. Indeed, if a path satisfies $q \,\mathsf{W}\, r$ then every prefix of the path also satisfies $q \,\mathsf{W}\, r$. We use proof by contradiction to argue that there is a path from t that does not satisfy $q \,\mathsf{W}\, r$ and does not pass through t'. Assume such a path does not exist. Then every path beginning in t that does not satisfy $q \,\mathsf{W}\, r$ has to pass through t'. However, both t and t' correspond to state s' in M. It follows that the only option to falsify $q \,\mathsf{W}\, r$ in game $\mathsf{G}_M(s', \alpha)$ is by "going in a loop" from state s' to itself. But by assumption all states on the path between t and t' satisfy q, a contradiction. □