

Regular Games with Imperfect Information Are Not That Regular

Laurent Doyen 

CNRS & LMF, ENS Paris-Saclay, France

Thomas Soullard 

LMF, ENS Paris-Saclay & CNRS, France

Abstract

We consider two-player games with imperfect information and the synthesis of a randomized strategy for one player that ensures the objective is satisfied almost-surely (i.e., with probability 1), regardless of the strategy of the other player. Imperfect information is modeled by an indistinguishability relation describing the pairs of histories that the first player cannot distinguish, a generalization of the traditional model with partial observations. The game is regular if it admits a regular function whose kernel commutes with the indistinguishability relation.

The synthesis of pure strategies that ensure all possible outcomes satisfy the objective is possible in regular games, by a generic reduction that holds for all objectives. While the solution for pure strategies extends to randomized strategies in the traditional model with partial observations (which is always regular), a similar reduction does not exist in the more general model. Despite that, we show that in regular games with Büchi objectives the synthesis problem is decidable for randomized strategies that ensure the outcome satisfies the objective almost-surely.

2012 ACM Subject Classification Theory of computation → Logic and verification; Theory of computation → Probabilistic computation

Keywords and phrases Imperfect-information games, randomized strategies, synthesis

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2024.1

1 Introduction

We consider the synthesis problem for two-player turn-based games with imperfect information, a model that has applications in several areas of computer-science, including the design of multi-agent systems [18], logics with uncertainty in planning and AI [5, 14], program synthesis [10], and automata theory [1]. The synthesis problem asks to decide the existence of (and if so, to construct) a strategy for one player that ensures a given objective is satisfied with the largest possible probability, regardless of the choices of the other player. We focus mostly on reachability objectives, and will also discuss ω -regular objectives [17].

Synthesis provides a natural formulation for the design of a reactive system that interacts with an unknown environment. The interactive nature of reactive systems is modeled by a two-player game between the system (the player) and an adversarial environment, and the limited access of the system to the current state of the game is modeled by imperfect information.

A simple example of a turn-based game with imperfect information is repeated matching pennies, where the environment (secretly) chooses head or tail (say, denoted by $x \in \{0, 1\}$), then the player chooses $y \in \{0, 1\}$ without seeing x , and wins if $y = x$. If $y \neq x$, the game repeats for one more round; the player loses if the game repeats forever. It is clear that there exists no pure (deterministic) strategy that is winning in repeated matching pennies, as for all sequences $\bar{x} = x_1 \dots x_i$ and $\bar{y} = y_1 \dots y_i$ that represent a history of the game (namely, the sequence of moves played in the first i rounds), given x_{i+1} (chosen by the environment), the choice of y_{i+1} (by the player) must be made independently of x_{i+1} , thus



© Laurent Doyen and Thomas Soullard;
licensed under Creative Commons License CC-BY 4.0

35th International Conference on Concurrency Theory (CONCUR 2024).

Editors: Rupak Majumdar and Alexandra Silva; Article No. 1; pp. 1:1–1:19

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

we may have $x_{i+1} = 1 - y_{i+1}$, showing that the game may repeat forever. That the player cannot ensure to win is counter-intuitive, and this is because pure (deterministic) strategies are considered. With randomized strategies, the player can win with probability 1, simply by choosing $y_i \in \{0, 1\}$ with uniform probability at every round. The power of randomization for decision making with imperfect information is well known and occurs in many situations [16, Section 1.3], including gene mutation in biology, penalty kick in sports, bluffing in card games, etc. It is also known that, even for simple objectives like reachability, optimal strategies may need memory [21].

In the sequel, we distinguish the synthesis of a pure strategy (which we call a sure winning strategy, as it aims to ensure all possible outcomes satisfy the objective), and the synthesis of a randomized strategy (which we call an almost-sure winning strategy, as it aims to ensure that the objective is satisfied with probability 1). Note that the synthesis problem for strategies that ensure the objective is satisfied with probability at least λ (given $0 < \lambda < 1$) is undecidable for probabilistic automata [20], a model that can be reduced to two-player turn-based games with imperfect information [11].

Given a finite alphabet Γ of moves, we denote by Γ^* the set of all histories in a game. The traditional model of imperfect information in games is a (regular) observation function defined on Γ^* that assigns to each history of the game a color that is visible to the player, while the history itself is not visible [22]. We call them partial-observation games, or games *à la Reif*. This model of games with imperfect information (even turn-based and non-stochastic) with randomized strategies generalizes many models such as concurrent games (e.g., matching pennies), Markov chains, Markov decision processes, and stochastic games [11, Theorem 5]. The synthesis of pure strategies in partial-observation games *à la Reif* with ω -regular objectives can be done using automata-based techniques analogous to the subset construction for finite automata [22, 21]. The synthesis of randomized (almost-sure winning) strategies can be done for reachability (and Büchi) objectives with a more involved technique based on the same approach of subset construction. The synthesis problem is EXPTIME-complete in both cases. Note that for ω -regular (and even for coBüchi) objectives, the synthesis problem with randomized strategies is undecidable [4].

A more abstract (and more general) model of imperfect information is given by a function $f : \Gamma^* \times \Gamma^* \rightarrow \{0, 1\}$ that specifies an indistinguishability relation [6, 7]: two histories τ, τ' are indistinguishable if $f(\tau, \tau') = 1$. In analogy to the partial-observation model, indistinguishable histories have the same length and their prefixes (of the same length) are also indistinguishable. Hence this function can be viewed as a set of pairs of finite words of the same length (over alphabet Γ), or alternatively as a language of finite words over the alphabet $\Gamma \times \Gamma$. Applications of this model include the synthesis of full-information protocols, which cannot be expressed by a partial-observation game *à la Reif* [7, Section 6].

A generic approach to the synthesis problem with this more general model is to construct a so-called rectangular morphism h defined on Γ^* with finite range [7]. Note that a morphism with finite range is a regular function (which can be defined by a finite-state automaton with output), and we say that a game is *regular* if such a rectangular morphism h exists. For instance, all partial-observation games *à la Reif* are regular as they admit a natural rectangular morphism, which is essentially the subset construction. The existence of a rectangular morphism guarantees that the information tree has a finite bisimulation quotient, and thus (pure) strategies can be transferred back and forth (by copying the action played in a bisimilar position of the other game), preserving the outcome of the game for arbitrary objectives [7]. Hence, for the synthesis of pure strategies, regular games with imperfect information can be reduced to finite games with perfect information.

In this paper, we consider the synthesis problem with randomized strategies, which is central in games with imperfect information (such as matching pennies). Given the existence of a finite bisimulation quotient in regular games, the fact that a solution based on a rectangular morphism (the subset construction) works for the synthesis of randomized strategies in reachability games *à la Reif* gives hope that synthesis in regular games is solvable along a similar path, by a reduction to a simpler equivalent game from which randomized strategies can be transferred back and forth, copying both the action and the probability of playing that action. The hope is reinforced by the existence of a bijection between the set of randomized strategies in a game *à la Reif* and the set of strategies in the simpler equivalent game, which relates strategies with same (probabilistic measure on the) outcome for arbitrary objectives [21, Theorem 4.2].

Surprisingly, we show that even a much weaker variant of such a reduction does not exist for regular games. First the hope for a bijective transfer of strategies is not reasonable because the information tree of regular games have in general unbounded branching [6], whereas the simpler equivalent games induced by (finite) rectangular morphism have a bounded-branching information tree. On the other hand, we show that even a non-bijective transfer of strategies, as general as it can reasonably be, may not exist for some regular game with a specific rectangular morphism (Section 4.3). Despite their nice structural properties, the existence of a finite bisimulation in their information tree, and their apparent similarity with games *à la Reif* where a bijection between randomized strategies is induced by a rectangular morphism, regular games are not as well-behaved as games *à la Reif*.

In this context, we present an algorithmic solution of the synthesis problem with randomized strategies for regular games with a reachability objective. The solution exploits the properties of rectangular morphisms to define a fixpoint computation with complexity quadratic in the size of the range of the rectangular morphism (which is of exponential size in the case of games *à la Reif* [21], and of non-elementary size in the case of full-information protocols [7]). Our algorithm shares the common features of the solutions of almost-sure reachability objectives in Markov decision processes [12], concurrent games [13], and games *à la Reif* [21], namely a nested fixpoint computation that iteratively constructs the almost-sure winning set by computing the set S_{i+1} of states from which the player can win with positive probability while ensuring to never leave the set S_i , where S_0 is the entire state space (Section 3). Our solution immediately extends to Büchi objectives, using reductions to reachability objectives from the literature.

In conclusion, this paper generalizes the positive (decidability) results about randomized strategies from games *à la Reif* to regular games (namely for reachability and Büchi objectives), whereas the reductions that worked in games *à la Reif* for arbitrary objectives no longer hold.

Omitted proofs are provided in the appendix.

2 Definitions

We recall basic definitions from logic and automata theory that will be useful in the rest of the paper, and then we discuss our model of games.

2.1 Basic notions

Regular functions. A *Moore machine* is a finite-state automaton with outputs, consisting of a finite input alphabet Γ , a finite output alphabet Σ , and a tuple $\mathcal{M} = \langle Q, q_\varepsilon, \delta, \lambda \rangle$ with a finite set Q of states (sometimes called the memory), an initial state $q_\varepsilon \in Q$, a (deterministic) transition function $\delta : Q \times \Gamma \rightarrow Q$, and an output function $\lambda : Q \rightarrow \Sigma$.

We extend the transition function to input words in the natural way, defining $\delta : Q \times \Gamma^* \rightarrow Q$ by $\delta(q, \varepsilon) = q$ for all states $q \in Q$, and inductively $\delta(q, \tau c) = \delta(\delta(q, \tau), c)$ for all histories $\tau \in \Gamma^*$ and moves $c \in \Gamma$.

The object of interest is the extension of the output function to a function $\lambda : \Gamma^* \rightarrow \Sigma$ defined by $\lambda(\tau) = \lambda(\delta(q_\varepsilon, \tau))$ for all histories $\tau \in \Gamma^*$. A function λ defined on Γ^* is *regular* if it is (the extension of) the output function of a Moore machine. The *cumulative extension* of λ is denoted by $\hat{\lambda} : \Gamma^* \rightarrow \Sigma^*$ and defined by $\hat{\lambda}(\varepsilon) = \varepsilon$ and inductively $\hat{\lambda}(\tau c) = \hat{\lambda}(\tau)\lambda(\tau c)$, thus $\hat{\lambda}(c_1 \dots c_k) = \lambda(c_1)\lambda(c_1 c_2) \dots \lambda(c_1 \dots c_k)$. Note that $\lambda = \text{last} \circ \hat{\lambda}$ where $\text{last}(c_1 \dots c_k) = c_k$. The function $\hat{\lambda}$ is further extended to infinite words as expected: $\hat{\lambda}(\pi) = \lambda(c_1)\lambda(c_1 c_2) \dots$ for all $\pi = c_1 c_2 \dots \in \Gamma^\omega$.

Fixpoint formulas. We briefly recall the interpretation of μ -calculus fixpoint formulas [8, 9] based on the Knaster-Tarski theorem. Given a monotonic function $\psi : 2^Q \rightarrow 2^Q$ (i.e., such that $X \subseteq Y$ implies $\psi(X) \subseteq \psi(Y)$), the expression $\nu Y. \psi(Y)$ is the (unique) greatest fixpoint of ψ , which can be computed as the limit of the sequence $(Y_i)_{i \in \mathbb{N}}$ defined by $Y_0 = Q$, and $Y_i = \psi(Y_{i-1})$ for all $i \geq 1$. Dually, the expression $\mu X. \psi(X)$ is the (unique) least fixpoint of ψ , and the limit of the sequence $(X_i)_{i \in \mathbb{N}}$ defined by $X_0 = \emptyset$, and $X_i = \psi(X_{i-1})$ for all $i \geq 1$. If $|Q| = n$, then it is not difficult to see that the limit of those sequences is reached after at most n iterations, $X_n = X_{n+1}$ and $Y_n = Y_{n+1}$.

On equivalence relations. Given an equivalence relation $\sim \subseteq S \times S$ and a set $T \subseteq S$, the *closure* of T under \sim is the set $[T]_\sim = \{t \in S \mid \exists t' \in T : t \sim t'\}$, and the *interior* of T under \sim is the set $\text{int}_\sim(T) = \{t \in S \mid [t]_\sim \subseteq T\}$. The closure and interior are monotone operators, that is if $T \subseteq U$, then $[T]_\sim \subseteq [U]_\sim$ and $\text{int}_\sim(T) \subseteq \text{int}_\sim(U)$. Note the duality $S \setminus \text{int}_\sim(T) = [S \setminus T]_\sim$, the complement of the interior of T is the closure of the complement of T . We say that T is *closed* under \sim if $T = [T]_\sim$ (or equivalently $T = \text{int}_\sim(T)$).

2.2 Games with imperfect information

We consider an abstract model of games, given by a (nonempty) set A of actions, a set Γ of moves, and a surjective function $\text{act} : \Gamma \rightarrow A$. The game is played over infinitely many rounds (or steps). In each round, the player chooses an action $a \in A$, then the environment chooses a move $c \in \Gamma$ such that $\text{act}(c) = a$. We say that the move c is *supported* by the action a .

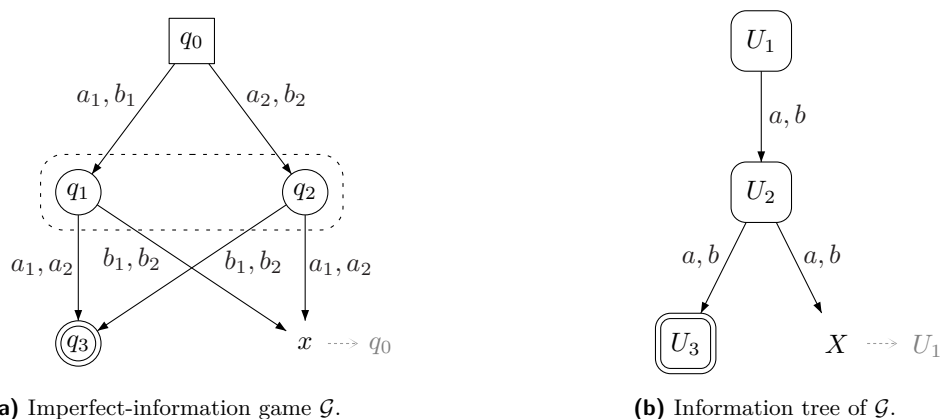
The outcome of the game is an infinite sequence $\pi \in \Gamma^\omega$ of moves of the environment, called a *play*, from which we can reconstruct the infinite sequence of actions of the player using the function act . A finite sequence of moves is called a *history*.

The winning condition is a set $W \subseteq \Gamma^\omega$ of plays that are winning for the player. We specify the winning condition by the combination of a coloring function $\lambda : \Gamma^* \rightarrow C$ for some finite alphabet C of colors, and a logical objective $L \subseteq C^\omega$. The coloring function is regular, and specified by a Moore machine. The induced winning condition is the set $W = \{\pi \in \Gamma^\omega \mid \hat{\lambda}(\pi) \in L\}$ of plays whose coloring satisfies the objective.

We consider the following classical objectives, for $C = \{0, 1\}$:

- the reachability objective is $\text{Reach} = 0^*1\{0, 1\}^\omega$ and for $x \in \mathbb{N}$ let $\text{Reach}^{\leq x} = \bigcup_{i \leq x} 0^i 1\{0, 1\}^\omega$;
- the safety objective is $\text{Safe} = 0^\omega = \{0, 1\}^\omega \setminus \text{Reach}$;
- the Büchi objective is $\text{Büchi} = (0^*1)^\omega$;
- the coBüchi objective is $\text{coBüchi} = \{0, 1\}^*0^\omega = \{0, 1\}^\omega \setminus \text{Büchi}$;

It is convenient to consider that the coloring function λ is part of a game instance, while the objective L can be specified independently of the game. This allows to quantify separately the game instances and the objectives. We recall that the traditional games



■ **Figure 1** A game with imperfect information representing matching pennies, and its information tree.

played on graphs [17, 3] are an equivalent model, in particular they can be translated into our abstract model of game, for example by letting moves be the edges of the game graph and the corresponding action be the label of the edge, while the graph structure can be encoded in the winning condition [23]. When we refer to a game played on a graph, we assume that all move sequences that contain two consecutive moves (i.e., edges) $c_i = (q_1, q_2)$ and $c_{i+1} = (q_3, q_4)$ such that $q_2 \neq q_3$ belong to the winning condition (for the player).

Imperfect information is represented by an indistinguishability relation $\sim \in \Gamma^* \times \Gamma^*$ that specifies the pairs of histories that the player cannot distinguish. An indistinguishability relation is an equivalence that satisfies the following conditions [6], for all histories $\tau, \tau' \in \Gamma^*$ and moves $c, c' \in \Gamma$: (i) if $\tau \sim \tau'$, then $|\tau| = |\tau'|$ (indistinguishable histories have the same length); (ii) if $\tau c \sim \tau' c'$, then $\tau \sim \tau'$ (the relation is prefix-closed) and $\text{act}(c) = \text{act}(c')$ (the action is visible). Indistinguishability relations may be specified using two-tape automata, but the results of this paper hold for arbitrary indistinguishability relations.

Given a history $\tau \in \Gamma^*$ we call the equivalence class $[\tau]_{\sim} = \{\tau' \in \Gamma^* \mid \tau' \sim \tau\}$ containing τ an *information set*. A function λ defined on Γ^* is *information consistent* if λ is constant over every information set: $\lambda(\tau) = \lambda(\tau')$ for all indistinguishable histories $\tau \sim \tau'$. An equivalent requirement is that the cumulative extension $\hat{\lambda}$ is constant over every information set, since indistinguishability relations are prefix-closed. We always require the coloring function of a game to be information consistent (with respect to the indistinguishability relation of the game). For every action $a \in A$, we define a successor relation \rightarrow_a over information sets, where $[\tau]_{\sim} \rightarrow_a [\tau c]_{\sim}$ for all moves $c \in \Gamma$ such that $\text{act}(c) = a$. The relations \rightarrow_a are obviously acyclic, and induce a structure called the *information tree*. Note that there is a natural extension of the coloring function λ to information sets, defined by $\lambda([\tau]_{\sim}) = \lambda(\tau)$ for all histories $\tau \in \Gamma^*$, since λ is information-consistent.

Given the above definitions, we represent a game with imperfect information as a tuple $\mathcal{G} = \langle A, \text{act}, \sim, \lambda \rangle$, together with an objective $L \subseteq C^\omega$, where the (finite) alphabet of moves is Γ , and the (finite) alphabet of colors is C . The information tree of \mathcal{G} is denoted by $\mathcal{U}_{\sim}(\mathcal{G})$ (consisting of the information sets of \mathcal{G} , the successor relations \rightarrow_a , and the coloring λ extended to information sets). The special case of perfect-information games corresponds the indistinguishability relation \sim being the identity (or equivalently by the information sets $[\tau]_{\sim} = \{\tau\}$ being singletons for all histories $\tau \in \Gamma^*$).

In figures, we present a very informal (but readable) description of games. The main components of a game that figures represent are the coloring function and the indistinguishability relation. Strictly speaking, the figures do not show a game but provide sufficient information to reconstruct the Moore machine defining the coloring function, and the two-tape automaton defining the indistinguishability relation. It should be possible to infer the color of a history (when its color is relevant), and the information set containing a given history. States are normally depicted as circles, but we may use boxes to emphasize that the action choice of the player is not relevant and only the move of the environment determines the successor state.

In the game \mathcal{G} of Figure 1a, representing matching pennies, the player has two actions, $A = \{a, b\}$, and the environment has two possible responses for each of them, thus four moves, $\Gamma = \{a_1, a_2, b_1, b_2\}$ where $\text{act}(a_i) = a$ and $\text{act}(b_i) = b$ for $i = 1, 2$. Intuitively, the adversarial environment chooses head (1) or tail (2) in the first round (at q_0) by choosing a response 1 or 2 (the action of the player is not relevant) that the player cannot observe (all histories of length 1 are indistinguishable, as suggested by the dashed line in Figure 1a). The player chooses head (action a) or tail (action b) in the next round, wins if the choice is matching the environment's choice, and replays the same game otherwise (x is a placeholder for the whole tree of Figure 1a).

The information tree of \mathcal{G} is shown in Figure 1b where $U_1 = [\varepsilon]_{\sim} = \{\varepsilon\}$ and $U_2 = [a_1]_{\sim} = \{a_1, a_2, b_1, b_2\} = \Gamma$. The information set U_3 contains eight of the sixteen histories of length 2. The node X is a placeholder for a tree that is isomorphic to the tree rooted at U_1 .

We recall that the nodes in information trees have a finite number of successors (there is at most $|\Gamma|^n$ information sets containing histories of length n), however the branching degree in a given information tree may be unbounded [6]. The partial-observation games *à la Reif* [22] are characterized by their information tree having bounded branching [6, Theorem 7].

2.3 Strategies and outcome

A *probability distribution* on a finite set S is a function $d : S \rightarrow [0, 1]$ such that $\sum_{s \in S} d(s) = 1$. A *Dirac distribution* assigns probability 1 to a (unique) element $s \in S$. We denote by $\mathcal{D}(S)$ the set of all probability distributions on S .

A strategy for the player is a function $\alpha : \Gamma^* \rightarrow \mathcal{D}(A)$ that is information consistent, thus $\alpha(\tau) = \alpha(\tau')$ for all indistinguishable histories $\tau \sim \tau'$. The strategy α is *pure* (or, deterministic) if $\alpha(\tau)$ is a Dirac distribution for all $\tau \in \Gamma^*$. A strategy for the environment is a function $\beta : \Gamma^* \times A \rightarrow \mathcal{D}(\Gamma)$ such that for all histories $\tau \in \Gamma^*$, actions $a \in A$, and moves $c \in \Gamma$, if the move c is played with positive probability $\beta(\tau, a)(c) > 0$, then it is supported by the action a , $\text{act}(c) = a$. The strategy β is *pure* (or, deterministic) if $\beta(\tau, a)$ is a Dirac distribution for all $\tau \in \Gamma^*$ and $a \in A$. We denote by $\mathcal{A}_{\mathcal{G}}$ (resp., by $\mathcal{B}_{\mathcal{G}}$) the set of all strategies for the player (resp., for the environment), and by $\mathcal{A}_{\mathcal{G}}^{\text{pure}}$ (resp., by $\mathcal{B}_{\mathcal{G}}^{\text{pure}}$) the set of all pure strategies for the player (resp., for the environment).

A pair of strategies (α, β) for the player and the environment induce a probability measure $\text{Pr}^{\alpha, \beta}$ on the sigma-algebra over the set of (infinite) plays. By Carathéodory's extension theorem, it is sufficient to define the probability measure over cylinder sets spanned by (finite) prefixes of plays. We define a family of probability measures $\text{Pr}_{\tau}^{\alpha, \beta}$ where $\tau \in \Gamma^*$ corresponds to starting the game after history τ . We set $\text{Pr}^{\alpha, \beta} = \text{Pr}_{\varepsilon}^{\alpha, \beta}$. Given $\rho \in \Gamma^*$, the measure of the cylinder set $\text{Cyl}(\rho) = \rho\Gamma^{\omega} = \{\pi \in \Gamma^{\omega} \mid \rho \text{ is a prefix of } \pi\}$ is defined as follows:

- if ρ is a prefix of τ , then

$$\text{Pr}_{\tau}^{\alpha, \beta}(\text{Cyl}(\rho)) = 1;$$

- if $\rho = \tau c_1 \dots c_k$ where $c_1, \dots, c_k \in \Gamma$ for some $k \geq 1$, then

$$\Pr_{\tau}^{\alpha, \beta}(\text{Cyl}(\rho)) = \prod_{i=1}^k \alpha(\tau c_1 \dots c_{i-1})(\text{act}(c_i)) \cdot \beta(\tau c_1 \dots c_{i-1}, \text{act}(c_i))(c_i);$$

- all other cylinder sets have measure 0.

Given an objective $L \subseteq C^\omega$ defined on colors, assuming the game \mathcal{G} (and thus $\hat{\lambda}$) is clear from the context we take the freedom to denote by $\Pr^{\alpha, \beta}(L)$ the probability $\Pr^{\alpha, \beta}(\{\pi \in \Gamma^\omega \mid \hat{\lambda}(\pi) \in L\})$ that the objective L is satisfied by an outcome of the pair of strategies (α, β) in \mathcal{G} . A strategy α of the player is *almost-sure winning* from an information set $[\tau]_{\sim}$ (or simply *almost-sure winning* if $\tau = \varepsilon$) for objective L if $\Pr_{\tau'}^{\alpha, \beta}(L) = 1$ for all $\tau' \in [\tau]_{\sim}$ and all strategies $\beta \in \mathcal{B}_{\mathcal{G}}$ of the environment. We are interested in solving the synthesis problem, which is to decide, given a game \mathcal{G} and objective L , whether there exists an almost-sure winning strategy in \mathcal{G} for L .

Note that the definition of almost-sure winning is equivalent to a formulation where only *pure* strategies of the environment are considered, that is where we require $\Pr_{\tau'}^{\alpha, \beta}(L) = 1$ for all $\beta \in \mathcal{B}_{\mathcal{G}}^{\text{pure}}$ only, as once a strategy for the player is fixed, the environment plays in a (possibly infinite-state) Markov decision process, for which pure strategies are sufficient [19, 11].

We say that a history $\tau \in \Gamma^*$ is *compatible* with α and β if $\Pr^{\alpha, \beta}(\text{Cyl}(\tau)) > 0$. An *outcome* of the pair of strategies (α, β) is a play $\pi \in \Gamma^\omega$ all of whose prefixes are compatible with α and β . We note that if $\Pr^{\alpha, \beta}(\text{Cyl}(\tau)) > 0$, then $\Pr_{\tau}^{\alpha, \beta}(\text{Cyl}(\tau c_1 c_2 \dots c_n)) = \Pr^{\alpha, \beta}(\text{Cyl}(\tau c_1 c_2 \dots c_n) \mid \text{Cyl}(\tau))$.

3 Almost-Sure Reachability

Throughout this section, we consider games with a (fixed) reachability objective $\text{Reach} = 0^*1\{0, 1\}^\omega$ over color alphabet $C = \{0, 1\}$, thus a play is winning if it has a finite prefix with color 1 according to the coloring function λ of the game. By extension, we say that a history τ is winning if $\lambda(\tau) = 1$. We also assume without loss of generality that if $\lambda(\tau) = 1$, then $\lambda(\tau c) = 1$ for all $\tau \in \Gamma^*$ and all $c \in \Gamma$. In the rest of this section, we fix a game \mathcal{G} as defined in Section 2.2 and we show how to decide the existence of an almost-sure winning strategy in games with a reachability objective, relying on the existence of a rectangular morphism (defined below). Two interesting classes of games are known to admit rectangular morphisms, namely the partial-observation games *à la Reif* and the full-information protocols [7].

3.1 Regular games with imperfect information

A *rectangular morphism* for \mathcal{G} is a surjective function $h : \Gamma^* \rightarrow P$ (for some finite set P) such that for all $\tau, \tau' \in \Gamma^*$ and $c \in \Gamma$:

Rectangularity¹ if $h(\tau) = h(\tau')$, then $h([\tau]_{\sim}) = h([\tau']_{\sim})$,

Morphism if $h(\tau) = h(\tau')$, then $h(\tau c) = h(\tau' c)$,

Refinement if $h(\tau) = h(\tau')$, then $\lambda(\tau) = \lambda(\tau')$.

Note that, as h is a finite-valued morphism, it can be represented by an automaton (whose output is not relevant) with input alphabet Γ and state space P . Games that admit a finite-valued rectangular morphism are called *regular*. Define the relation $\approx = \{(h(\tau), h(\tau')) \mid \tau \sim \tau'\}$ and recall a non-trivial fundamental result.

¹ Rectangularity is equivalent to the kernel $H = \{(\tau, \tau') \mid h(\tau) = h(\tau')\}$ of h commuting with the indistinguishability relation, that is $H \circ \sim = \sim \circ H$.

► **Lemma 1.** [7, Lemma 3 & Lemma 4] *The relation \approx is an equivalence and $h([\tau]_{\sim}) = [h(\tau)]_{\approx}$ for all $\tau \in \Gamma^*$.*

We call the elements of P *abstract states*. We extend the relation \approx to elements of $P \times A$ as follows: for all $p, p' \in P$ and $a, a' \in A$, let $(p, a) \approx (p', a')$ if $p \approx p'$ and $a = a'$.

Thanks to the morphism property, there exists a function $\delta^P : P \times \Gamma \rightarrow P$ such that $\delta^P(h(\tau), c) = h(\tau c)$ for all $\tau \in \Gamma^*$ and $c \in \Gamma$. Define the set $P_F = \{p \in P \mid \exists \tau \in \Gamma^* : h(\tau) = p \wedge \lambda(\tau) = 1\}$ of target (or final) states, which are the images by h of winning histories. Note that the coloring function is information-consistent and therefore the set P_F is closed under \approx by Lemma 1. Moreover, by our assumption on the coloring function the states in P_F form a sink set, that is $\delta^P(p, c) \in P_F$ for all states $p \in P_F$ and moves $c \in \Gamma$.

Rectangular morphisms are central to the solution of the synthesis problem for *pure* strategies [7], by showing that the following abstract game is equivalent to \mathcal{G} : starting from $p_0 = h(\varepsilon)$, the game is played in rounds where each round starts with a value p and the player chooses an action $a \in A$, the next round starts in p' chosen by the environment such that $p' \approx \delta^P(p, c)$ for some move $c \in \Gamma$ such that $\text{act}(c) = a$. The player wins if a value in P_F eventually occurs along a play. We may view the elements of P as positions of the player, and elements of $P \times A$ as positions of the environment. We further discuss this abstract game in Section 4.2.

An abstract state $p \in P$ is *existentially winning* if there exists a history $\tau \in \Gamma^*$ such that $h(\tau) = p$ and the player has an *almost-sure winning* strategy from $[\tau]_{\sim}$. We denote by P_{\exists} the set of all existentially winning states. Note that the set P_{\exists} is closed under \approx , since for all $p \in P_{\exists}$ and $p' \approx p$, by Lemma 1 there exists $\tau' \in [\tau]_{\sim}$, such that $h(\tau') = p'$ and the player has an *almost-sure winning* strategy from $[\tau']_{\sim} = [\tau]_{\sim}$.

An abstract state $p \in P$ is *universally winning* if for all histories $\tau \in \Gamma^*$ such that $h(\tau) = p$, the player has an *almost-sure winning* strategy from $[\tau]_{\sim}$. We denote by P_{\forall} the set of all universally winning states. Note that the set P_{\forall} is closed under \approx , by an argument analogous to the case of P_{\exists} . As we consider a reachability objective, it is easy to see that $P_F \subseteq P_{\forall} \subseteq P_{\exists}$ (the latter inclusion holds since h is surjective).

Note that in games of perfect information, the existentially and universally winning states coincide by definition since information sets are singletons, $P_{\forall} = P_{\exists}$. A corollary of our results is that $P_{\forall} = P_{\exists}$ also holds in games of imperfect information with a reachability objective.

3.2 Algorithm

We present an algorithm to decide if the player is almost-sure winning for reachability in a game of imperfect information, assuming we have a rectangular morphism.

Given a set $X \subseteq P \cup (P \times A)$, define the *predecessor* operator as follows:

$$\begin{aligned} \text{Pre}(X) &= \{p \in P \mid \exists a \in A : (p, a) \in X\} \\ &\cup \\ &\{(p, a) \in P \times A \mid \forall c : \text{act}(c) = a \implies \delta^P(p, c) \in X\} \end{aligned}$$

It is easy to verify that if $X_1 \subseteq X_2$, then $\text{Pre}(X_1) \subseteq \text{Pre}(X_2)$, that is $\text{Pre}(\cdot)$ is monotone. Intuitively, from a history τ such that $h(\tau) \in \text{Pre}(X)$ the player can choose an action a such that $(h(\tau), a) \in X$, and if $(h(\tau), a) \in \text{Pre}(X)$ then for all moves c chosen by the environment and supported by action a we have $h(\tau c) \in X$. In a game of perfect information, iterating the

predecessor operator from the target set P_F until obtaining a fixpoint $X_* = \mu X. \text{Pre}(X) \cup P_F$ gives the (existentially and universally) winning states, that is $P_{\forall} = P_{\exists} = X_*$.

In a game of imperfect information with randomized strategies for the player, given $X_{i+1} = \text{Pre}(X_i)$, from a history τ such that $h(\tau) \in X_{i+1}$, we may consider playing all actions a such that $(h(\tau), a) \in X_i$ uniformly at random. However, an action played in a history τ is also played in every indistinguishable history $\tau' \sim \tau$. Therefore, we need to ensure that for all actions a played in τ (such that $(h(\tau), a) \in X_i$), playing a in τ' does not leave X_* . Hence for $p = h(\tau)$, even if $(p, a) \in X_*$, the action a should not be played from τ if there exists $p' \approx p$ such that $(p', a) \notin X_*$.

In our algorithm, we remove from X_* all elements (p, a) such that $(p', a) \notin X_*$ for some $p' \approx p$, that is we replace X_* by its interior. We define

$$Y_* = \nu Y. \mu X. \text{int}_{\approx}(Y) \cap (\text{Pre}(X) \cup P_F)$$

and we show that the sets Y_* , P_{\forall} , and P_{\exists} coincide. Note that the μ -calculus formula for Y_* is well defined since $\text{Pre}(\cdot)$ and $\text{int}_{\approx}(\cdot)$ are monotone operators, and that the fixpoint can be effectively computed since P is finite.

As the fixpoint satisfies $Y_* = \text{int}_{\approx}(Y_*) \cap (\text{Pre}(Y_*) \cup P_F)$, it follows that $Y_* = \text{int}_{\approx}(Y_*)$ is closed under \approx , and $Y_* \subseteq \text{Pre}(Y_*) \cup P_F$. Moreover it is easy to verify from the fixpoint iteration that $P_F \subseteq Y_*$ since P_F , which is closed under \approx , is always contained in the least fixpoint, and never removed from the greatest fixpoint.

► **Theorem 2.** *The abstract states in the fixpoint Y_* are the existentially winning states, which coincide with the universally winning states, $Y_* \cap P = P_{\forall} = P_{\exists}$.*

Theorem 2 entails the decidability of the existence of an almost-sure winning strategy in games with reachability objective: if there exists an almost-sure winning strategy (from ε), then $h(\varepsilon) \in P_{\exists} = Y_*$, and conversely if $h(\varepsilon) \in Y_*$ then $h(\varepsilon) \in P_{\forall}$ since $h(\varepsilon) \in P$, and there exists an almost-sure winning strategy from $[\varepsilon]_{\sim} = \{\varepsilon\}$.

► **Corollary 3.** *There exists an almost-sure winning strategy if and only if $h(\varepsilon) \in Y_*$.*

The result of Theorem 2 follows from the chain of inclusions $Y_* \cap P \subseteq P_{\forall}$ (Lemma 4), the already established $P_{\forall} \subseteq P_{\exists}$, and $P_{\exists} \subseteq Y_* \cap P$ (Lemma 5),

To show that $Y_* \cap P \subseteq P_{\forall}$ we construct a strategy for the player that is almost-sure winning from all τ (and from $[\tau]_{\sim}$) such that $h(\tau) \in Y_*$. The strategy plays uniformly at random all actions that ensure the successor τc of τ remains in Y_* , more precisely $h(\tau c) \in Y_*$. We show that at least one such action ensures progress towards reaching a target state in P_F , thus with probability at least $\frac{1}{|A|}$. The target is reachable in at most $|P|$ steps, which entails bounded probability (at least $\nu = \frac{1}{|A|^{|P|}}$) to reach P_F (from every state of Y_*) and since the strategy ensures that Y_* is never left, the probability of *not* reaching P_F is at most $(1 - \nu)^k$ for all $k \geq 0$, and as $(1 - \nu)^k \rightarrow 0$ when $k \rightarrow \infty$ the probability to eventually reach P_F is 1.

► **Lemma 4.** $Y_* \cap P \subseteq P_{\forall}$.

The last inclusion of the chain is proved by showing that there exists a fixpoint that contains P_{\exists} and since Y_* is defined as the greatest fixpoint, we have $P_{\exists} \subseteq Y_*$.

► **Lemma 5.** $P_{\exists} \subseteq Y_* \cap P$.

The fixpoint computation for Y_* can be implemented by a quadratic algorithm, with respect to the number $|P|$ of abstract states (assuming a constant number of actions).

Both stochastic games on graphs and Büchi objectives are subsumed by the results of this paper: in stochastic games, almost-sure Büchi reduces to almost-sure reachability [4, Lemma 8.3], and games with stochastic transitions can be simulated by (non-stochastic) games with imperfect information [11, Theorem 5]. These results can easily be lifted to the more general framework of games with imperfect information defined by indistinguishability relations. On the other hand, we recall that almost-sure coBüchi is undecidable, already in probabilistic automata with pure strategies [4], and also with randomized strategies [11, Corollary 2], but without the assumption that the coloring function is information consistent. For the special case of information-consistent coBüchi objectives, the decidability status of the synthesis problem for almost-sure winning is open (even in games *à la Reif*), to the best of our knowledge. Finally for safety objectives, almost-sure winning is equivalent to sure winning (which requires that all possible outcomes satisfy the objective), and the problem is equivalent to synthesis with pure strategies [7].

► **Theorem 6.** *Given a regular game \mathcal{G} with imperfect information and a rectangular morphism $h : \Gamma^* \rightarrow P$ for \mathcal{G} , the synthesis problem for almost-sure reachability and Büchi objectives can be solved in time $O(|P|^2)$.*

Since $|P|$ is already exponential in games *à la Reif* (for which the synthesis problem is EXPTIME-complete), the quadratic blow-up is not significant. For example, we get a $(k+1)$ -EXPTIME complexity upper bound for the synthesis problem in the model of full-information protocols (FIP) with k observers, using the rectangular morphism of [7, Section 5.2]. The model of full-information protocols as presented in [7] features explicit communication actions that entail full disclosure of all available information of the sender. The game involves several players who may communicate, but only one active player who takes control actions, which reduces the synthesis problem to a game with imperfect information as considered here.

A matching $(k+1)$ -EXPTIME lower bound can be obtained by a straightforward adaptation of the proof of [7, Theorem 3], which presents a reduction from the membership problem for alternating k -EXPSpace Turing machines to the synthesis problem of (pure) winning strategies. The reduction is such that in the constructed game, randomization does not help the player. If the player takes a chance in deviating from the faithful simulation of the Turing machine, a losing sink state is reached, thus with positive probability.

► **Theorem 7.** *The synthesis problem for FIP games with k observers (and almost-sure reachability objective) is $(k+1)$ -EXPTIME-complete.*

The construction of an almost-sure winning strategy α can be done by following the first step in the proof of Lemma 4, which constructs α given the value of the fixpoint Y_* . It follows that the constructed strategy is a regular function that can be represented by the automaton underlying the rectangular morphism of the game, and thus memory of size $|P|$ is sufficient to define an almost-sure winning strategy.

4 Reductions

The results of Section 3 may suggest the existence of a (strong) correspondence, possibly an equivalence, between the game \mathcal{G} with regular indistinguishability relation and a simpler (finite-state) game \mathcal{H} (essentially the abstract game presented in Section 3.1) that would hold for arbitrary objectives.

While it is virtually impossible to establish the non-existence of a reasonably strong such correspondence, we show for a natural notion of game equivalence that such a correspondence does not hold in general.

4.1 Alternating probabilistic trace equivalence

Inspired by the notion of alternating refinement relations [2] in non-stochastic game graphs with perfect information (also called alternating transition systems), we consider the following definition of *game refinement*.

Given two games \mathcal{G} and \mathcal{H} , we say that \mathcal{G} *reduces* to \mathcal{H} (denoted by $\mathcal{G} \preceq \mathcal{H}$) if for all strategies $\alpha_{\mathcal{G}}$ of the player in \mathcal{G} , there exists a strategy $\alpha_{\mathcal{H}}$ of the player in \mathcal{H} , such that for all strategies $\beta_{\mathcal{H}}$ of the environment in \mathcal{H} , there exists a strategy $\beta_{\mathcal{G}}$ of the environment in \mathcal{G} , such that the probability measures $\Pr^{\alpha_{\mathcal{G}}, \beta_{\mathcal{G}}}$ in \mathcal{G} and $\Pr^{\alpha_{\mathcal{H}}, \beta_{\mathcal{H}}}$ in \mathcal{H} coincide (on all objectives $L \subseteq C^\omega$, that is $\Pr^{\alpha_{\mathcal{G}}, \beta_{\mathcal{G}}}(L) = \Pr^{\alpha_{\mathcal{H}}, \beta_{\mathcal{H}}}(L)$ for all L), in short:

$$\forall \alpha_{\mathcal{G}} \cdot \exists \alpha_{\mathcal{H}} \cdot \forall \beta_{\mathcal{H}} \cdot \exists \beta_{\mathcal{G}} \cdot \forall L \subseteq C^\omega : \Pr^{\alpha_{\mathcal{G}}, \beta_{\mathcal{G}}}(L) = \Pr^{\alpha_{\mathcal{H}}, \beta_{\mathcal{H}}}(L).$$

As expected, the quantifications over strategies of the player range over information-consistent randomized strategies (in their respective game), and the quantification over L ranges over measurable sets. Under these quantifications, the condition $\Pr^{\alpha_{\mathcal{G}}, \beta_{\mathcal{G}}}(L) = \Pr^{\alpha_{\mathcal{H}}, \beta_{\mathcal{H}}}(L)$ can be replaced by $\Pr^{\alpha_{\mathcal{G}}, \beta_{\mathcal{G}}}(L) \leq \Pr^{\alpha_{\mathcal{H}}, \beta_{\mathcal{H}}}(L)$ (consider the complement of L , which is also a measurable set), and then game refinement can be interpreted as the game \mathcal{H} being easier² than \mathcal{G} for the player because for every strategy of the player in \mathcal{G} there is a better² strategy for the player in \mathcal{H} (as reflected by the quantifier sequence $\forall \alpha_{\mathcal{G}} \cdot \exists \alpha_{\mathcal{H}}$), which means that the environment can always ensure a lower² value in \mathcal{G} (against $\alpha_{\mathcal{G}}$) than in \mathcal{H} (against $\alpha_{\mathcal{H}}$), as reflected by the quantifier sequence $\forall \beta_{\mathcal{H}} \cdot \exists \beta_{\mathcal{G}}$.

Note that in the special case of pure strategies (and non-stochastic games) where the probability measures assign probability 1 to a single play, game refinement boils down to

$$\forall \alpha_{\mathcal{G}} \in \mathcal{A}_{\mathcal{G}}^{\text{pure}} \cdot \exists \alpha_{\mathcal{H}} \in \mathcal{B}_{\mathcal{H}}^{\text{pure}} : \text{outcome}_{\mathcal{H}}(\alpha_{\mathcal{H}}) \subseteq \text{outcome}_{\mathcal{G}}(\alpha_{\mathcal{G}})$$

where $\text{outcome}_X(\alpha)$ denotes the set of plays compatible with α in game X . The condition $\text{outcome}_{\mathcal{H}}(\alpha_{\mathcal{H}}) \subseteq \text{outcome}_{\mathcal{G}}(\alpha_{\mathcal{G}})$ is equivalent to

$$\forall \beta_{\mathcal{H}} \in \mathcal{B}_{\mathcal{H}}^{\text{pure}} \cdot \exists \beta_{\mathcal{G}} \in \mathcal{B}_{\mathcal{G}}^{\text{pure}} : \text{outcome}_{\mathcal{H}}(\alpha_{\mathcal{H}}, \beta_{\mathcal{H}}) = \text{outcome}_{\mathcal{G}}(\alpha_{\mathcal{G}}, \beta_{\mathcal{G}}),$$

where $\text{outcome}_X(\alpha, \beta)$ denotes the set of plays compatible with α and β in game X . This is essentially the definition of alternating trace containment [2, Section 4]. Hence our definition of game refinement is a natural generalization of both alternating trace containment to a stochastic and imperfect-information setting, and of refinement for labeled Markov decision processes [15, Section 4] to games.

We say that two games \mathcal{G} and \mathcal{H} are *equivalent* (or inter-reducible) if $\mathcal{G} \preceq \mathcal{H}$ and $\mathcal{H} \preceq \mathcal{G}$. We sometimes refer to this equivalence as *alternating probabilistic trace equivalence*.

4.2 Bijection and bisimulation

We compare the approach in Section 3 and the technique for solving almost-sure reachability in games with partial observation *à la Reif* [21]. Recall that games with partial observation are a special case of games with regular indistinguishability relation.

The solution of the synthesis problem with randomized strategies for games *à la Reif* established a bijection (denoted by h , but we call it \hat{h} to emphasize that it is a mapping between sequences) between the histories in the original game \mathcal{G} and in a game \mathcal{H} of partial

² The words “easier”, “better”, “lower”, etc. are interpreted in a non-strict sense.

observation played as follows: given a history ρ (initially ε), the player chooses an action $a \in A$, and the environment extends the history to $\rho' = \hat{h}(\tau c)$ where $\hat{h}(\tau) = \rho$ and c is a move (in \mathcal{G}) such that $\text{act}(c) = a$. Two histories are indistinguishable if all their respective prefixes (of the same length) are \approx -equivalent (according to the definition of \approx before Lemma 1). We call \mathcal{H} the *abstract game induced* from \mathcal{G} by \hat{h} , and by an abuse of notation, we denote by \approx the indistinguishability relation of \mathcal{H} .

The bijection \hat{h} naturally extends to a bijection between plays in \mathcal{G} and plays in \mathcal{H} , and further to a bijection between strategies in \mathcal{G} and strategies in \mathcal{H} that preserves probability measures. The existence of this bijection immediately ensures that the games \mathcal{G} and \mathcal{H} are equivalent (in the sense of alternating probabilistic trace equivalence). Another important consequence is that the information trees of \mathcal{G} and \mathcal{H} are isomorphic, a situation that is not guaranteed when \mathcal{G} is a game with regular indistinguishability relation and \mathcal{H} is a finite-state game (i.e., the set $h(\Gamma^*) = \{\text{last}(\hat{h}(\tau)) \mid \tau \in \Gamma^*\}$ is finite). Indeed, the branching degree of the information tree of \mathcal{G} may be unbounded [6], whereas the branching degree of the information tree of \mathcal{H} is necessarily bounded since \mathcal{H} is a finite game. It is therefore impossible to follow the same route as in games *à la Reif*, using bijections.


Since bijection is too strong for our general setting, a more realistic hope is to rely on the existence of a bisimulation between the information tree of \mathcal{G} and the information tree of \mathcal{H} [7, Theorem 1]. Note that bisimulation is insensitive to the branching degree. The bisimulation ensures that, under pure strategies for the player, the (perfect-information) games played on the information trees are equivalent (in the sense of alternating trace equivalence). The equivalence induced by the bisimulation is such that the transfer of strategies (the construction of $\alpha_{\mathcal{H}}$ given $\alpha_{\mathcal{G}}$) is of the form $\alpha_{\mathcal{H}}(\cdot) = \alpha_{\mathcal{G}}(\mu(\cdot))$ where μ is a mapping between bisimilar histories³ that is sequential, that is the image of a prefix of a history is a prefix of the image of that history. So, for pure strategies the action played by $\alpha_{\mathcal{H}}$ at a history is a copy of the action played by $\alpha_{\mathcal{G}}$ in a bisimilar history in \mathcal{G} . In the case of randomized strategies, the player chooses a distribution over actions, which we may view as attaching probabilities to actions, in a way that could be copied along with the action to transfer randomized strategies (on one hand from \mathcal{G} to \mathcal{H} , and on the other hand from \mathcal{H} to \mathcal{G}) and establish alternating probabilistic trace equivalence (containment in both directions) of the two games.

We recall the strong relationships between games with imperfect information and their induced abstract game (Figure 2). First, there is a natural bijection between an imperfect-information game \mathcal{G} and the perfect-information game played on its information tree $\mathcal{U}_{\sim}(\mathcal{G})$, which is a witness of alternating trace equivalence of the two games [7, Lemma 1]. Second, assuming the existence of a rectangular morphism h for \mathcal{G} , the information trees of \mathcal{G} and of \mathcal{H} (the abstract game induced from \mathcal{G} by h) are bisimilar. The bijections and the bisimulation (Figure 2) entail alternating trace equivalence of the games \mathcal{G} and \mathcal{H} and their information trees, and thus equivalence of the synthesis problems with pure strategies in \mathcal{G} and in \mathcal{H} [7, Lemma 2, Theorem 1].

4.3 Counterexamples

We present two examples showing the difficulty to adapt the results of alternating trace equivalence for games with pure strategies, to alternating probabilistic trace equivalence for games with randomized strategies.

³ In games played over information trees, the histories are sequences of information sets; we do not denote them by a symbol to avoid having to define them formally.



■ **Figure 2** Equivalences between the game \mathcal{G} , its induced abstract game \mathcal{H} , and their information trees.

First, we show that with randomized strategies, playing on the information tree of a game with imperfect information is not equivalent to playing in the original game, although this is true for pure strategies (links (1) of Figure 2).

The first example is matching pennies (the game \mathcal{G} in Figure 1a). Recall that the set of actions is $A = \{a, b\}$, the set of moves is $\Gamma = \{a_1, a_2, b_1, b_2\}$ with $\text{act}(a_i) = a$ and $\text{act}(b_i) = b$ for $i = 1, 2$. It is well known that the player cannot win matching pennies using a pure strategy, but wins almost-surely by choosing head and tail uniformly at random in every round. However, in the game played on the information tree of \mathcal{G} (Figure 1b), the player cannot win, even with positive probability. After the first round the information set is $U_2 = \Gamma$, which has two successors on both actions a and b . This is because U_2 contains both a history τ such that τa_i is winning and a history τ' such that $\tau' b_i$ is winning, hence U_3 is a successor of U_2 (on both actions a and b), and analogously X is a successor of U_2 (X is a placeholder for the whole tree of Figure 1b). Therefore, after any choice of action by the player in U_2 , the environment can always choose X and continue the game, ensuring that the reachability objective is never satisfied.

This first example illustrates a crucial difference between pure and randomized strategies. Considering a game \mathcal{G} (with imperfect information induced by an indistinguishability relation \sim) and the game played on its information tree $\mathcal{U}_{\sim}(\mathcal{G})$ (which is of perfect information), there is a natural bijection between (information-consistent) strategies in \mathcal{G} and strategies in $\mathcal{U}_{\sim}(\mathcal{G})$. The bijection exists for both pure and randomized strategies. However, only in the case of pure strategies, the bijection is a witness of alternating trace equivalence (Figure 2). Intuitively, this is because once a *pure* strategy for the player is fixed, the environment “knows” the specific action played in every history (which is the same within an information set), and therefore constructing a path in the original game \mathcal{G} compatible with the player strategy is equivalent to constructing a path in $\mathcal{U}_{\sim}(\mathcal{G})$ (informally, the information tree is a valid abstraction for “existence of paths”). In contrast, once a *randomized* strategy for the player is fixed, the environment only “knows” the specific distribution over actions played in every history, which leaves an element of surprise as to which action will be played when the strategy is executed. In this context, the information tree is no longer a valid abstraction, intuitively because it allows the environment to choose a history within the current information set *after* the randomly chosen action is known, as illustrated by matching pennies (Figure 1b). To be valid, the abstraction should force the environment to stick to a single choice of a history, before the action is drawn from the distribution chosen by the player. We may consider the abstract game \mathcal{H} induced from \mathcal{G} as a candidate for a richer abstraction, and establish a direct link (2) between \mathcal{G} and \mathcal{H} (Figure 2), given the

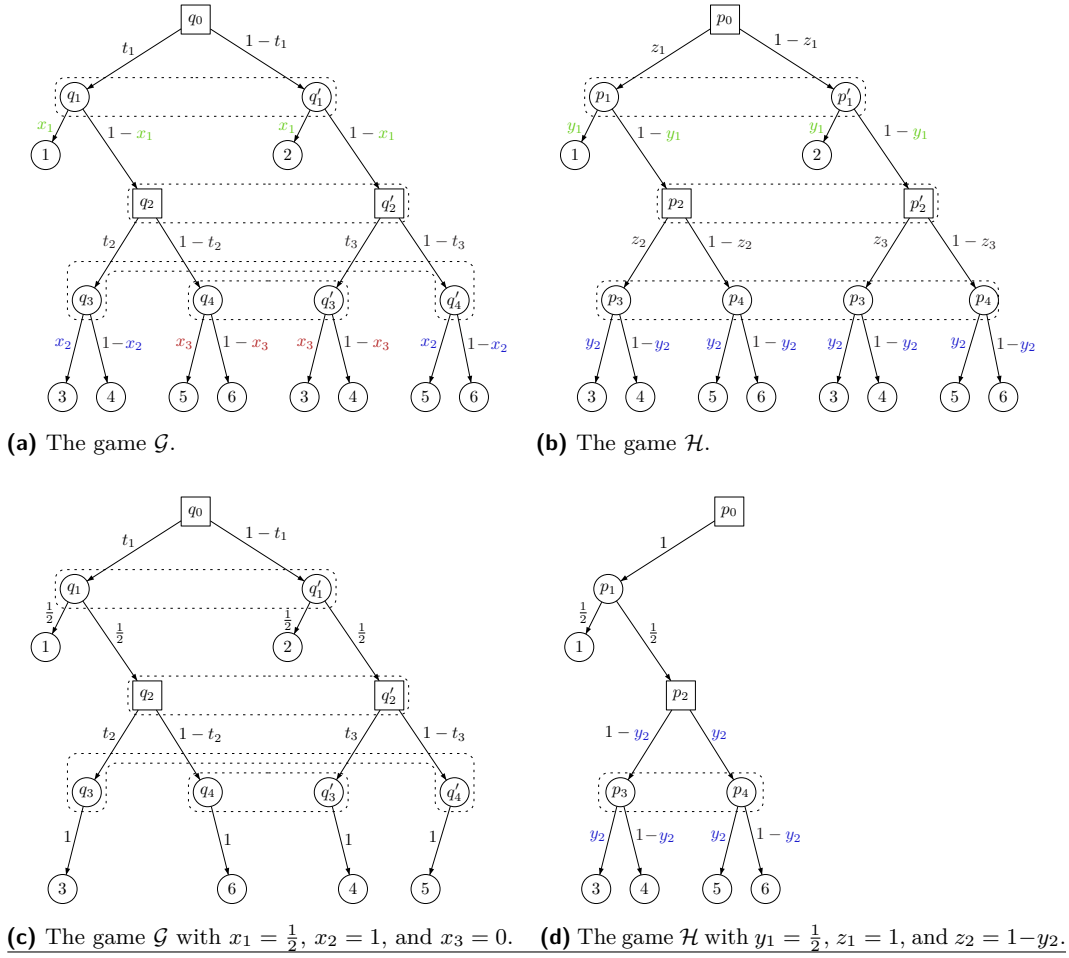


Figure 3 A game \mathcal{G} and the abstract game \mathcal{H} induced by a rectangular morphism, with randomized strategies encoded by the variables \bar{x}, \bar{t} (in \mathcal{G}) and \bar{y}, \bar{z} (in \mathcal{H}).

link (1) breaks. Unfortunately, in our second example \mathcal{H} is not equivalent to \mathcal{G} for alternating probabilistic trace equivalence ($\mathcal{G} \not\sim \mathcal{H}$), showing the absence of such a direct link, which however exists in partial-observation games *à la Reif*, via the bijection \hat{h} mentioned in Section 4.2.

The second example (Figure 3a) shows a game \mathcal{G} and its abstract game \mathcal{H} induced by a rectangular morphism. To avoid tedious description of the game, we present the key features informally. The game \mathcal{G} has two actions for the player and two possible responses for the adversary. The actions and moves are not shown in Figure 3a, and states are depicted as circles when only the action of the player determines the successor, and as boxes when only the move of the environment determines the successor. All states have color 0 except the leaves, labeled by their color 1, 2, ..., 6. The leaves are sink states (with a self-loop on every move). The dashed lines show the indistinguishability classes. It is easy to check that the coloring function is information-consistent.

Figure 3b shows the abstract game \mathcal{H} , as a tree with the same shape as in Figure 3a for \mathcal{G} . The value $h(\tau)$ of a history can be read as the label of the node in Figure 3b corresponding

to the path defined by τ in Figure 3a. For example, the histories (of length 3) in \mathcal{G} that correspond to the nodes q_3 and q'_3 are mapped by h to the abstract state p_3 . Figure 3b can be viewed as the unraveling of the morphism h . It is easy to verify that h is a rectangular refinement of the coloring function, and that \approx -equivalent abstract states are grouped within the dashed lines.

We show that \mathcal{G} does not reduce to \mathcal{H} according to alternating probabilistic trace containment ($\mathcal{G} \not\leq \mathcal{H}$). We describe all randomized strategies of the player and of the environment by specifying the probability to play an action or a move in each (relevant, i.e., non-leaf) history. We use variables $\bar{x} = x_1, x_2, x_3$ for $\alpha_{\mathcal{G}}$; $\bar{y} = y_1, y_2$ for $\alpha_{\mathcal{H}}$; $\bar{z} = z_1, z_2, z_3$ for $\beta_{\mathcal{H}}$; and $\bar{t} = t_1, t_2, t_3$ for $\beta_{\mathcal{G}}$. Note that the player may only use information-consistent strategies, hence the same variable is used from indistinguishable histories to describe $\alpha_{\mathcal{G}}$ and $\alpha_{\mathcal{H}}$. The claim $\mathcal{G} \not\leq \mathcal{H}$ can be written as

$$\psi \equiv \exists \bar{x} \cdot \forall \bar{y} \cdot \exists \bar{z} \cdot \forall \bar{t} : \bigvee_{c=1}^6 \varphi_c$$

where φ_c expresses that the mass of probability on color c differs in \mathcal{G} and in \mathcal{H} . For example,

$$\varphi_1 \equiv t_1 x_1 \neq z_1 y_1, \text{ and}$$

$$\varphi_3 \equiv t_1 t_2 x_2 (1 - x_1) + t_3 x_3 (1 - t_1) (1 - x_1) \neq z_1 z_2 y_2 (1 - y_1) + z_3 y_2 (1 - z_1) (1 - y_1).$$

To show that ψ holds, we fix $x_1 = \frac{1}{2}$, $x_2 = 1$, and $x_3 = 0$. Now, consider all possible values of \bar{y} and observe that if $y_1 \neq \frac{1}{2}$, then ψ holds since the probability mass in colors $\{1, 2\}$ is $\frac{1}{2}$ in \mathcal{G} and y_1 in \mathcal{H} (and thus $\varphi_1 \vee \varphi_2$ must hold regardless of the value of \bar{z} and \bar{t}). So, it remains to show that ψ holds for $y_1 = \frac{1}{2}$ (and for all values of y_2). Fix $z_1 = 1$ and $z_2 = 1 - y_2$ (the value of z_3 is arbitrary), as illustrated in Figure 3c and Figure 3d.

Consider all possible values of \bar{t} and observe that if $t_1 \neq 1$, then ψ holds since the probability mass in color 1 is $\frac{1}{2}$ in \mathcal{G} and t_1 in \mathcal{H} (thus φ_1 holds). With $t_1 = 1$, we show that $\varphi_4 \vee \varphi_5$ holds, that is:

$$0 \neq \frac{(1 - y_2)^2}{2} \vee 0 \neq \frac{(y_2)^2}{2}$$

which holds since either $y_2 \neq 0$ or $1 - y_2 \neq 0$.

We conclude that ψ holds and thus \mathcal{G} does not reduce to \mathcal{H} ($\mathcal{G} \not\leq \mathcal{H}$). Note that slightly stronger statements can be proved by a similar argument: first, the strategy $\beta_{\mathcal{H}}$ (represented by the variables \bar{z}) can be chosen pure, by setting $z_2 = 1$ if $y_2 = 0$ and $z_2 = 0$ if $y_2 = 1$ (and setting z_2 to an arbitrary value in $\{0, 1\}$ otherwise); second, the example does not even preserve almost-sure probabilities, in the following sense. Given a target set $T \subseteq C = \{1, \dots, 6\}$ of colors, let $\text{Reach}(T) = \bigcup_{k \in T} C^* k C^\omega$ be the reachability objective with target T . Consider the formula

$$\psi' \equiv \exists \bar{x} \cdot \forall \bar{y} \cdot \exists \bar{z} \cdot \forall \bar{t} : \bigvee_{T \subseteq C} \neg(\text{Pr}_{\mathcal{G}}^{\bar{x}, \bar{t}}(\text{Reach}(T)) = 1) \Leftrightarrow \text{Pr}_{\mathcal{H}}^{\bar{y}, \bar{z}}(\text{Reach}(T)) = 1)$$

where $\text{Pr}_{\mathcal{G}}^{\bar{x}, \bar{t}}(\text{Reach}(T))$ is the probability that the objective $\text{Reach}(T)$ is satisfied in \mathcal{G} under strategies $\alpha_{\mathcal{G}}$ defined by the variables \bar{x} for the player and $\beta_{\mathcal{G}}$ defined by \bar{t} for the environment (and analogously for $\text{Pr}_{\mathcal{H}}^{\bar{y}, \bar{z}}(\text{Reach}(T))$ in \mathcal{H}). It is easy to check that the condition ψ' entails ψ .

To prove that ψ' holds, we fix $x_1 = \frac{1}{2}$, $x_2 = 1$, and $x_3 = 0$ as before (Figure 3c). Consider three possible cases for \bar{y} : (i) if $y_1 = 0$, then no matter the value of the other variables the objective $\text{Reach}(\{1, 2\})$ has probability $\frac{1}{2}$ in \mathcal{G} and probability 0 in \mathcal{H} , so we take $T = C \setminus \{1, 2\} = \{3, 4, 5, 6\}$; (ii) if $y_1 = 1$, we take $T = \{1, 2\}$ (the probability mass is $\frac{1}{2}$

in \mathcal{G} and 1 in \mathcal{H}); (iii) otherwise $0 < y_1 < 1$, and we take $z_1 = 1$ and $z_2 = 1 - y_2$ (illustrated in Figure 3d for $y_1 = \frac{1}{2}$), and consider all possible values of \bar{t} : if $t_1 \neq 1$ then $\text{Reach}(\{2\})$ has positive probability in \mathcal{G} and probability 0 in \mathcal{H} , so we take $T = C \setminus \{2\} = \{1, 3, 4, 5, 6\}$, and if $t_1 = 1$, then $\text{Reach}(\{4, 5\})$ has probability 0 in \mathcal{G} and positive probability in \mathcal{H} , so we take $T = C \setminus \{4, 5\} = \{1, 2, 3, 6\}$.

► **Theorem 8.** *There exists a regular game \mathcal{G} such that the abstract game \mathcal{H} induced from \mathcal{G} is not alternating probabilistic trace equivalent to \mathcal{G} .*

References

- 1 Shaull Almagor, Udi Boker, and Orna Kupferman. What’s decidable about weighted automata? *Inf. Comput.*, 282:104651, 2022. doi:10.1016/j.ic.2020.104651.
- 2 Rajeev Alur, Thomas A. Henzinger, Orna Kupferman, and Moshe Y. Vardi. Alternating refinement relations. In Davide Sangiorgi and Robert de Simone, editors, *CONCUR ’98: Concurrency Theory, 9th International Conference, Nice, France, September 8-11, 1998, Proceedings*, volume 1466 of *Lecture Notes in Computer Science*, pages 163–178. Springer, 1998. doi:10.1007/BFb0055622.
- 3 Krzysztof R. Apt and Erich Grädel, editors. *Lectures in Game Theory for Computer Scientists*. Cambridge University Press, 2011.
- 4 Christel Baier, Marcus Größer, and Nathalie Bertrand. Probabilistic ω -automata. *J. ACM*, 59(1):1, 2012. doi:10.1145/2108242.2108243.
- 5 Chitta Baral, Thomas Bolander, Hans van Ditmarsch, and Sheila A. McIlraith. Epistemic planning (Dagstuhl seminar 17231). *Dagstuhl Reports*, 7(6):1–47, 2017. doi:10.4230/DagRep.7.6.1.
- 6 Dietmar Berwanger and Laurent Doyen. Observation and distinction. representing information in infinite games. In Christophe Paul and Markus Bläser, editors, *37th International Symposium on Theoretical Aspects of Computer Science, STACS 2020, March 10-13, 2020, Montpellier, France*, volume 154 of *LIPICs*, pages 48:1–48:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.STACS.2020.48.
- 7 Dietmar Berwanger, Laurent Doyen, and Thomas Soullard. Synthesising full-information protocols. *CoRR*, abs/2307.01063, 2023. arXiv:2307.01063, doi:10.48550/arXiv.2307.01063.
- 8 Julian C. Bradfield and Colin Stirling. Modal mu-calculi. In Patrick Blackburn, J. F. A. K. van Benthem, and Frank Wolter, editors, *Handbook of Modal Logic*, volume 3 of *Studies in logic and practical reasoning*, pages 721–756. North-Holland, 2007. doi:10.1016/s1570-2464(07)80015-2.
- 9 Julian C. Bradfield and Igor Walukiewicz. The mu-calculus and model checking. In Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, editors, *Handbook of Model Checking*, pages 871–919. Springer, 2018. URL: https://doi.org/10.1007/978-3-319-10575-8_26, doi:10.1007/978-3-319-10575-8_26.
- 10 Pavol Cerný, Krishnendu Chatterjee, Thomas A. Henzinger, Arjun Radhakrishna, and Rohit Singh. Quantitative synthesis for concurrent programs. In Ganesh Gopalakrishnan and Shaz Qadeer, editors, *Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings*, volume 6806 of *Lecture Notes in Computer Science*, pages 243–259. Springer, 2011. URL: https://doi.org/10.1007/978-3-642-22110-1_20, doi:10.1007/978-3-642-22110-1_20.
- 11 Krishnendu Chatterjee, Laurent Doyen, Hugo Gimbert, and Thomas A. Henzinger. Randomness for free. *Inf. Comput.*, 245:3–16, 2015. doi:10.1016/j.ic.2015.06.003.
- 12 Costas Courcoubetis and Mihalis Yannakakis. The complexity of probabilistic verification. *J. ACM*, 42(4):857–907, 1995. doi:10.1145/210332.210339.
- 13 Luca de Alfaro, Thomas A. Henzinger, and Orna Kupferman. Concurrent reachability games. *Theor. Comput. Sci.*, 386(3):188–217, 2007. doi:10.1016/j.tcs.2007.07.008.

- 14 Stéphane Demri and Raul Fervari. Model-checking for ability-based logics with constrained plans. In Brian Williams, Yiling Chen, and Jennifer Neville, editors, *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, pages 6305–6312. AAAI Press, 2023. doi:10.1609/aaai.v37i15.25776.
- 15 Laurent Doyen, Thomas A. Henzinger, and Jean-François Raskin. Equivalence of labeled markov chains. *Int. J. Found. Comput. Sci.*, 19(3):549–563, 2008. doi:10.1142/S0129054108005814.
- 16 Robert Gibbons. *A primer in game theory*. Harvester Wheatsheaf, 1992.
- 17 Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002. doi:10.1007/3-540-36387-4.
- 18 Dilian Gurov, Valentin Goranko, and Edvin Lundberg. Knowledge-based strategies for multi-agent teams playing against nature. *Artif. Intell.*, 309:103728, 2022. doi:10.1016/j.artint.2022.103728.
- 19 Donald A. Martin. The determinacy of blackwell games. *J. Symb. Log.*, 63(4):1565–1581, 1998. doi:10.2307/2586667.
- 20 A. Paz. *Introduction to probabilistic automata*. Academic Press, 1971.
- 21 Jean-François Raskin, Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. Algorithms for omega-regular games with imperfect information. *Log. Methods Comput. Sci.*, 3(3), 2007. doi:10.2168/LMCS-3(3:4)2007.
- 22 John H. Reif. The complexity of two-player games of incomplete information. *J. Comput. Syst. Sci.*, 29(2):274–301, 1984. doi:10.1016/0022-0000(84)90034-5.
- 23 Wolfgang Thomas. On the synthesis of strategies in infinite games. In Ernst W. Mayr and Claude Puech, editors, *STACS 95, 12th Annual Symposium on Theoretical Aspects of Computer Science, Munich, Germany, March 2-4, 1995, Proceedings*, volume 900 of *Lecture Notes in Computer Science*, pages 1–13. Springer, 1995. URL: https://doi.org/10.1007/3-540-59042-0_57, doi:10.1007/3-540-59042-0_57.

A Proof of Lemma 4

The abstract states computed by the fixpoint Y_* are all universally winning.

► **Lemma 4.** $Y_* \cap P \subseteq P_V$.

Proof. Given $p \in Y_* \cap P$, we show that $p \in P_V$. Let $A_p = \{a \in A \mid (p, a) \in Y_*\}$ and show that $A_p \neq \emptyset$. Since $Y_* \subseteq \text{Pre}(Y_*) \cup P_F$, either (i) $p \in P_F$ and then $A_p = A \neq \emptyset$ as $P_F \times A \subseteq \text{Pre}(P_F)$ (states in P_F form a sink set and P_F is closed under \approx); or (ii) $p \in \text{Pre}(Y_*)$ and by definition there exists $a \in A$ such that $(p, a) \in Y_*$, hence $A_p \neq \emptyset$.

Consider the strategy α for the player defined, for all $\tau \in \Gamma^*$, by $\alpha(\tau) = d_U(A_p)$ where $p = h(\tau)$ and $d_U(\cdot)$ is the uniform distribution. We show that α is information-consistent. Given $\tau' \sim \tau$, let $p = h(\tau)$ and $p' = h(\tau')$. Hence we have $p \approx p'$, and since Y_* is closed under \approx , if $(p, a) \in Y_*$ then $(p', a) \in Y_*$. It follows that $A_p = A_{p'}$ and thus $\alpha(\tau) = \alpha(\tau')$.

The strategy α ensures that Y_* is never left: if $h(\tau) \in Y_*$ and $a \in \text{Supp}(\alpha(\tau))$, then $h(\tau c) \in Y_*$ for all c such that $\text{act}(c) = a$. Indeed, for $p = h(\tau)$ we have $a \in A_p$ and $(p, a) \in Y_*$ thus $(p, a) \in \text{Pre}(Y_*)$ which imply $\delta^P(p, c) = h(\tau c) \in Y_*$.

We now show that α ensures from Y_* a target state is reached within the next $N = |P|$ steps with probability at least $\nu = \frac{1}{|A|^N}$, that is $\Pr_\tau^{\alpha, \beta}(\text{Reach}^{\leq |\tau|+N}) \geq \nu$ for all strategies β of the environment and for all $\tau \in \Gamma^*$ such that $h(\tau) \in Y_*$.

Since $Y_* = \mu X. \text{int}_{\approx}(Y_*) \cap (\text{Pre}(X) \cup P_F)$, the set Y_* is the limit of the (non-decreasing) sequence $(X_i)_{i \in \mathbb{N}}$ where $X_0 = \emptyset$ and $X_{i+1} = Y_* \cap (\text{Pre}(X_i) \cup P_F)$. The *rank* of an element $y \in Y_*$ is the least $i \geq 0$ such that $y \in X_i$ (i.e., such that $y \in X_i$ and $y \notin X_{i-1}$). By extension, the rank of a history τ such that $h(\tau) \in Y_*$ is the rank of $h(\tau)$. Note that states $p \in Y_* \cap P$ have odd rank, and the state-action pairs $(p, a) \in Y_* \cap (P \times A)$ have even rank. The smallest rank is 1 and corresponds to the set P_F of target states. Let N_* be the largest rank, which is bounded by $|Y_*|$.

Intuitively, from any history τ such that $h(\tau) \in Y_*$ the strategy α ensures, against all strategies of the environment, that the history τc after the next round has a lower rank (unless the rank of τ is 1, and thus τ is a winning history) with probability at least $\frac{1}{|A|}$. Let $r > 1$ be the rank of $h(\tau)$, then $h(\tau) \in \text{Pre}(X_{r-1})$ and thus there is an action $a \in A$ such that $(h(\tau), a) \in X_{r-1}$, thus $(h(\tau), a) \in \text{Pre}(X_{r-2})$ and $h(\tau c) \in X_{r-2}$ for all c with $\text{act}(c) = a$. It follows that $a \in A_p$ is played with probability at least $\frac{1}{|A|}$ and the intuitive claim follows.

Combined with the fact that the strategy α ensures Y_* is never left, this claim shows that from all τ such that $h(\tau) \in Y_*$, against all strategies β of the environment, the reachability objective is satisfied (i.e., the rank decreases to 1) within the next N_* rounds with probability at least $\nu = \frac{1}{|A|^{N_*}}$. Formally $\Pr_\tau^{\alpha, \beta}(\text{Reach}^{\leq |\tau|+N_*}) \geq \nu$ against all strategies. It follows that the probability of *not* reaching a target state within $k \cdot N_*$ rounds is at most $(1 - \nu)^k$ which tends to 0 as $k \rightarrow \infty$, thus $\Pr_\tau^{\alpha, \beta}(\text{Reach}) = 1$ by the squeeze theorem. We conclude that the player is almost-sure winning from $[\tau]_{\sim}$ since Y_* is closed under \approx , which implies that $p = h(\tau) \in P_V$. ◀

B Proof of Lemma 5

The fixpoint Y_* contains all existentially winning states.

► **Lemma 5.** $P_{\exists} \subseteq Y_* \cap P$.

Proof. As a preliminary, it will be useful to note that if a strategy α_{as} for the player is almost-sure winning from an information set $[\tau]_{\sim}$, then for all actions $a \in \text{Supp}(\alpha_{\text{as}}(\tau))$, for all moves c supported by a , $\text{act}(c) = a$, the strategy α_{as} is almost-sure winning from $[\tau c]_{\sim}$. An equivalent conclusion is that for all actions $a \in \text{Supp}(\alpha_{\text{as}}(\tau))$, the pair $(h(\tau), a)$ is in $\text{int}_{\approx}(\text{Pre}(P_{\exists}))$ (\star).

We now proceed with the proof of the lemma. First, $P_{\exists} \subseteq P$ by definition. To show that $P_{\exists} \subseteq Y_*$, since Y_* is defined as a greatest fixpoint, it is sufficient to show that $Y_{\exists} := P_{\exists} \cup \text{int}_{\approx}(\text{Pre}(P_{\exists}))$ is a fixpoint, that is $Y_{\exists} = \mu X. \text{int}_{\approx}(Y_{\exists}) \cap (\text{Pre}(X) \cup P_F)$. Since P_{\exists} is closed under \approx (i.e., $P_{\exists} = \text{int}_{\approx}(P_{\exists})$) so is Y_{\exists} , and we only need to show that $Y_{\exists} \subseteq \mu X. Y_{\exists} \cap (\text{Pre}(X) \cup P_F) =: X_*$ (the converse inclusion is trivial).

Given $y \in Y_{\exists}$, we show that $y \in X_*$, which concludes the proof. We consider two cases, either $y = p \in P_{\exists}$ or $y = (p, a) \in \text{int}_{\approx}(\text{Pre}(P_{\exists}))$:

- (i) if $p \in P_{\exists}$, then there exists $\tau_p \in \Gamma^*$ such that $h(\tau_p) = p$ and the player is almost-sure winning from $[\tau_p]_{\sim}$ using a strategy α_{as} . Assume towards contradiction that $p \notin X_*$.

We construct a (pure) spoiling strategy β for the environment as follows. For all $\tau \in \Gamma^*$ such that $h(\tau) \in P_{\exists}$ and $h(\tau) \notin X_*$ (thus $h(\tau) \notin \text{Pre}(X_*)$), for all actions $a \in \text{Supp}(\alpha_{\text{as}}(\tau))$ played by α_{as} in τ , the pair $(h(\tau), a)$ is not in X_* , and further not in $\text{Pre}(X_*)$ since $(h(\tau), a) \in \text{int}_{\approx}(\text{Pre}(P_{\exists}))$ by (\star), thus there exists a move c supported by action a such that $h(\tau c) = \delta^P(h(\tau), c) \notin X_*$. Define $\beta(\tau, a) = c$, and define $\beta(\tau', a)$ arbitrarily for τ' such that $h(\tau') \in X_*$.

We have the following, for all histories $\tau \in \Gamma^*$ and for $a \in \text{Supp}(\alpha_{\text{as}}(\tau))$ and $c = \beta(\tau, a)$:

- if $h(\tau) \in P_{\exists}$, then $h(\tau c) \in P_{\exists}$ (since α_{as} is almost-sure winning),
- if $h(\tau) \in P_{\exists}$ and moreover $h(\tau) \notin X_*$, then $h(\tau c) \notin X_*$,
- $h(\tau_p) \in P_{\exists}$ and $h(\tau_p) \notin X_*$.

It follows by an inductive argument that against α , the constructed strategy β ensures from τ_p that X_* and thus P_F is never reached, hence $\Pr_{\tau_p}^{\alpha_{\text{as}}, \beta}(\text{Reach}) = 0$, which contradicts that α_{as} is almost-sure winning from τ_p . Hence $p \in X_*$.

- (ii) if $(p, a) \in \text{int}_{\approx}(\text{Pre}(P_{\exists}))$, then $\delta^P(p, c) \in P_{\exists}$ for all moves c supported by a . By an argument similar to case (i), we have $\delta^P(p, c) \in X_*$. It follows that $(p, a) \in \text{Pre}(X_*)$ and thus $(p, a) \in X_* = Y_{\exists} \cap (\text{Pre}(X_*) \cup P_F)$.

◀