

Admissibility in Concurrent Games

Nicolas Basset¹, Gilles Geeraerts¹, Jean-François Raskin¹, and
Ocan Sankur²

- 1 Université libre de Bruxelles, Brussels, Belgium
{nicolas.basset, gilles.geeraerts, jraskin}@ulb.ac.be
2 CNRS, IRISA, Rennes, France
ocan.sankur@irisa.fr

Abstract

In this paper, we study the notion of *admissibility* for randomised strategies in *concurrent games*. Intuitively, an admissible strategy is one where the player plays ‘as well as possible’, because there is no other strategy that *dominates* it, i.e., that wins (almost surely) against a super set of adversarial strategies. We prove that admissible strategies always exist in concurrent games, and we characterise them precisely. Then, when the objectives of the players are ω -regular, we show how to perform *assume-admissible synthesis*, i.e., how to compute admissible strategies that win (almost surely) under the hypothesis that the other players play admissible strategies only.

Digital Object Identifier 10.4230/LIPIcs...


1 Introduction

In a concurrent n -player game played on a graph, all n players *independently* and *simultaneously* choose moves at each round of the game, and those n choices determine the next state of the game [13]. Concurrent games generalise turn-based games and it is well-known that, while deterministic strategies are sufficient in the turn-based case, randomised strategies are necessary for winning with probability one even for reachability objectives. Intuitively, randomisation is necessary because in concurrent games, in each round, players have no information about the concurrent choice of moves made by the other players. Randomisation allows for some probability of choosing a good move while not knowing the choice of the other players. As a consequence, there are two classical semantics that are considered to analyse these games *qualitatively*: winning with certainty (sure semantics in the terminology of [13]), and winning with probability one (almost sure semantics in the terminology of [13]). We consider both semantics here.

Previous papers on concurrent games are mostly concerned with two-player zero-sum games, i.e. two players that have fully antagonistic objectives. In this paper, we consider the *more general setting* of n -player non zero-sum concurrent games in which each player has its own objective. The notion of winning strategy is not sufficient to study non zero-sum games and other solution concepts have been proposed. One such concept is the notion of *admissible strategy* [1].

For a player with objective Φ , a strategy σ is said to be *dominated* by a strategy σ' if σ' does as well as σ with respect to Φ against all the strategies of the other players and strictly better for some of them. A strategy σ is *admissible* for a player if it is not dominated by any other of his strategies. Clearly, playing a strategy which is not admissible is sub-optimal and a rational player should only play admissible strategies. While recent works have studied the notion of admissibility for n -player non zero-sum game graphs [4, 14, 9, 7, 6], they are all concerned with the special case of turn-based games and this work is the first to consider the more general concurrent games.






■ **Figure 1** The relationships between the classes of Admissible, LA, and SCO strategies for three families of games. All the inclusions are strict.

Throughout the paper, we consider the running example in Figure 2. This is a concurrent game played by two players. Player 1’s objective is to reach Trg , while Player 2 wants to reach s_2 . Edges are labelled by pairs of moves of both players which activate that transition (where $-$ means ‘any move’). It is easy to see that no player can enforce its objective with or without randomisation, so, there is no winning strategy in this game for either player. This is because moving from s_0 to s_1 and from s_1 to s_2 requires the cooperation of both players. Moreover, the transitions from s_2 behave as in the classical ‘matching pennies’ game: player 1 must choose between f and f' ; player 2 between g and g' ; and the target is reached only when the choices ‘match’. So, randomisation is needed to make sure Trg is reached with probability one, from s_2 . In the paper, we will describe the dominated and admissible strategies of this game.

Technical contributions First, we study the notion of admissible strategies for both the *sure* and *almost sure* semantics of concurrent games. We show in Theorem 8 that in both semantics admissible strategies always exist. The situation is thus similar to the turn-based case [4, 9]. Nevertheless, the techniques used in this simpler case do not generalise easily to the concurrent case and we need substantially more involved technical tools here. To obtain our universal existence result, we introduce two weaker solution concepts: *locally admissible moves* and *strongly cooperative optimal* strategies. While cooperative optimal strategies were already introduced in [6] and shown equivalent to admissible strategies in the turn based setting, there are strictly weaker than admissible strategies in the concurrent setting (both for the sure and the almost sure semantics), and they need to be combined with the notion of locally admissible moves to fully characterise admissible strategies. In the special case of safety objectives, we can show that admissible strategies are exactly those that always play locally admissible moves. This situation is depicted in Figure 1.

Second, we build on our characterisation of admissible strategies based on the notions of locally admissible moves and strongly cooperative optimal strategies to obtain algorithms to solve the *assume admissible synthesis* problem for concurrent games. In the assume admissible synthesis problem, we ask whether a given player has an *admissible* strategy that is *winning* against *all admissible* strategies of the other players. So this rule relaxes the classical synthesis rule by asking for a strategy that is winning against the admissible strategies of the other players only and not against all of them. This is reasonable as in a multi-player game, each player has his own objective which is generally not the complement of the objectives of the other players. The assume-admissible rule makes the hypothesis that players are rational, hence they play admissible strategies and it is sufficient to win against those strategies. Our algorithm is applicable to all ω -regular objectives and it is based on a reduction to a zero-sum two-player game in the sure semantics. While this reduction shares intuitions with the reduction that we proposed in [7] to solve the same problem in the



■ **Figure 2** A concurrent game where Player 1 and 2 want to reach Trg and s_2 respectively.

turn-based case, our reduction here is based on games with imperfect information [17]. In contrast, in the turn-based case, games of perfect information are sufficient. The correctness and completeness of our reduction are proved in Theorem 11.

Related works Concurrent reachability games were studied in [13] and algorithms to solve more general omega-regular objectives were given in [10]. The games studied there are two-player and zero-sum only. We rely on the algorithms defined in [10] to compute states from which players have almost surely winning strategies for their objectives when all the other players play adversely. States where a player has a (deterministic) strategy to surely win against all other players can be computed by a reduction to more classical turn-based game graphs [2]. Nash equilibria have been studied in concurrent games [5], but without randomised strategies. None of those papers consider the notion of admissibility.

We use the notion of admissibility to obtain synthesis algorithms for systems composed of several sub-systems starting from non zero-sum specifications. There have been a few other proposals in the literature that are based on refinements of the notion of Nash equilibrium (and not on admissibility), most notably: assume-guarantee synthesis [11] and rational synthesis [15, 16]. Those works assume the simpler setting of turn-based games and so they do not deal with randomised strategies. In the context of infinite games played on graphs, one well known limitation of Nash equilibria is the existence of non-credible threats, admissibility does not suffer from this problem.

In [12], Damm and Finkbeiner use the notion of *dominant strategy* to provide a compositional semi-algorithm for the (undecidable) distributed synthesis problem. So while we use the notion of admissible strategy, they use a notion of dominant strategy. The notion of dominant strategy is *strictly stronger*: every dominant strategy is admissible but an admissible strategy is not necessary dominant. Also, in multiplayer games with omega-regular objectives with complete information (as considered here), admissible strategies are always guaranteed to exist [4] while it is not the case for dominant strategies.

2 Preliminaries

Concurrent games played on graphs Let $P = \{1, 2, \dots, n\}$ be a set of players. A *concurrent game* played on a finite graph by the players in P is a tuple $\mathcal{G} = (S, \Sigma, s_{\text{init}}, (\Sigma_p)_{p \in P}, \delta)$ where, (i) S is a finite set of *states*; and $s_{\text{init}} \in S$ the *initial state*; (ii) Σ is a finite set of *actions*; (iii) For all $p \in P$, $\Sigma_p : S \rightarrow 2^\Sigma \setminus \{\emptyset\}$ is a *action assignment* that assigns, to all states $s \in S$, the set of actions available to player p from state s . (iv) $\delta : S \times \Sigma \times \dots \times \Sigma \rightarrow S$ is the *transition function*. We write $\Sigma(s) = \Sigma_1(s) \times \dots \times \Sigma_n(s)$ for all $s \in S$. It is often convenient to consider a player p separately and see the set of all other players $P \setminus \{p\}$ as a single player denoted $-p$. Hence, the set of actions of $-p$ in state s is: $\Sigma_{-p}(s) =_{\text{def}} \prod_{q \in P \setminus \{p\}} \Sigma_q(s)$. We assume that $\Sigma_i(s) \cap \Sigma_j(s) = \emptyset$ for all $s \in S$ and $i \neq j$. We denote by $\text{Succ}(s, a) = \{\delta(s, a, b) \mid b \in \Sigma_{-p}(s)\}$ the set of possible *successors* of the state $s \in S$ when player p performs action $a \in \Sigma_p(s)$. A particular case of concurrent games are

the *turn-based games*. A game $\mathcal{G} = (S, \Sigma, s_{\text{init}}, (\Sigma_p)_{p \in P}, \delta)$ is turn-based iff for all states $s \in S$, there is a unique player p s.t. the successors of s depend only on p 's choice of action, i.e., $\text{Succ}(s, a)$ contains exactly one state for all $a \in \Sigma_p(s)$.

A *history* is a finite path $h = s_1 s_2 \dots s_k \in S^*$ s.t. (i) $k \in \mathbb{N}$; (ii) $s_1 = s_{\text{init}}$; and (iii) for every $2 \leq i \leq k$, there exists $(a_1, \dots, a_n) \in \Sigma^{|P|}$ with $s_i = \delta(s_{i-1}, a_1, \dots, a_n)$. The *length* $|h|$ of a history $h = s_1 s_2 \dots s_k$ is its number of states k ; for every $1 \leq i \leq k$, we denote by h_i the state s_i and by $h_{\leq i}$ the history $s_1 s_2 \dots s_i$. We denote by $\text{last}(h)$ the last state of h , that is, $\text{last}(h) = h_{|h|}$. A *run* is defined similarly as a history except that its length is infinite. For a run $\rho = s_1 s_2 \dots \in S^\omega$ and $i \in \mathbb{N}$, we also write $\rho_{\leq i} = s_1 s_2 \dots s_i$ and $\rho_i = s_i$. Let $\text{Hist}(\mathcal{G})$ (resp. $\text{Runs}(\mathcal{G})$) denote the set of histories (resp. runs) of \mathcal{G} . The game is played from the initial state s_{init} for an infinite number of rounds, producing a run. At each round $i \geq 0$, with current state s_i , all players p select simultaneously a action $a_p^i \in \Sigma_p(s_i)$, and the state $\delta(s_i, a_1^i, \dots, a_n^i)$ is appended to the current history. The selection of the action by a player is done according to strategies defined below.

Randomised moves and strategies Given a finite set A , a probability distribution on A , is a function $\alpha : A \rightarrow [0, 1]$ such that $\sum_{a \in A} \alpha(a) = 1$; and we let $\text{Supp}(\alpha) = \{a \mid \alpha(a) > 0\}$ be the support of α . We denote by $\alpha(B) = \sum_{a \in B} \alpha(a)$ the probability of a given set B according to α . The set of probability distributions on A is denoted by $\mathcal{D}(A)$. A *randomised move* of player p in state s is a probability distribution on $\Sigma_p(s)$, that is, an element of $\mathcal{D}(\Sigma_p(s))$. A randomised move that assigns probability 1 to an action and 0 to the others is called a *Dirac move*. We will henceforth denote randomised moves as sums of actions weighted by their respective probabilities. For instance $0.5f + 0.5g$ denotes the randomised move that assigns probability 0.5 to f and g (and 0 to all other actions). In particular, we denote by b a Dirac move that assigns probability 1 to action b .

Given a state s and a tuple $\beta = (\beta_p)_{p \in P} \in \prod_{p \in P} \mathcal{D}(\Sigma_p(s))$ of randomised moves from s , one per player, we let $\delta_r(s, \beta) \in \mathcal{D}(S)$ be the probability distribution on states s.t. for all $s' \in S$: $\delta_r(s, \beta)(s') = \sum_{\mathbf{a} \mid \delta(s, \mathbf{a}) = s'} \beta(\mathbf{a})$, where $\beta(a_1, \dots, a_n) = \prod_{i=1}^n \beta_i(a_i)$. Intuitively, $\delta_r(s, \beta)(s')$ is the probability to reach s' from s when the players play according to β .

A *strategy* for player p is a function σ from histories to randomised moves (of player p) such that, for all $h \in \text{Hist}(\mathcal{G})$: $\sigma(h) \in \mathcal{D}(\Sigma_p(\text{last}(h)))$. A strategy is called Dirac at history h , if $\sigma(h)$ is a Dirac move; it is called Dirac if it is Dirac in all histories. We denote by $\Gamma_p(\mathcal{G})$ the set of player- p strategies in the game, and by $\Gamma_p^{\text{det}}(\mathcal{G})$ the set of player- p strategies that only use Dirac moves (those strategies are also called *deterministic*); we might omit \mathcal{G} if it is clear from context. A *strategy profile* σ for a subset $A \subseteq P$ of players is a tuple $(\sigma_p)_{p \in A}$ with $\sigma_p \in \Gamma_p$ for all $p \in A$. When the set of players A is omitted, we assume $A = P$. Let $\sigma = (\sigma_p)_{p \in P}$ be a strategy profile. Then, for all players p , we let σ_{-p} denote the restriction of σ to $P \setminus \{p\}$ (hence, σ_{-p} can be regarded as a strategy of player $-p$ that returns, for all histories h , a randomised move from $\prod_{p \in P \setminus \{p\}} \mathcal{D}(\Sigma_p(s)) \subseteq \mathcal{D}(\Sigma_{-p}(\text{last}(h)))$). We sometimes denote σ by the pair (σ_p, σ_{-p}) . Given a history h , we let $(\sigma_p)_{p \in A}(h) = (\sigma_p(h))_{p \in A}$.

Let h be a history and let ρ be a history or a run. Then, we write $h \subseteq_{\text{pref}} \rho$ iff h is a prefix of ρ , i.e., $\rho_{\leq |h|} = h$. Consider two strategies σ and σ' for player p , and a history h . We denote by $\sigma \langle h \leftarrow \sigma' \rangle$ the strategy that follows strategy σ and *shifts* to σ' as soon as h has been played. Formally, $\sigma \langle h \leftarrow \sigma' \rangle$ is the strategy s.t., for all histories h' : $\sigma \langle h \leftarrow \sigma' \rangle(h') = \sigma'(h')$ if $h \subseteq_{\text{pref}} h'$; and $\sigma \langle h \leftarrow \sigma' \rangle(h') = \sigma(h')$ otherwise.

Probability measure and outcome of a profile Given a history h , we let $\text{Cyl}(h) = \{\rho \mid h \subseteq_{\text{pref}} \rho\}$ be the *cylinder* of h . To each strategy profile σ , we associate a probability measure

\mathbb{P}_σ on certain sets of runs. First, for a history h , we define $\mathbb{P}_\sigma(\text{Cyl}(h))$ inductively on the length of h : $\mathbb{P}_\sigma(\text{Cyl}(s_{\text{init}})) = 1$, and $\mathbb{P}_\sigma(\text{Cyl}(h's')) = \mathbb{P}_\sigma(\text{Cyl}(h')) \cdot \delta_r(\text{last}(h'), \sigma(h'))(s')$ when $|h| > 1$ and $h = h's'$. Based on this definition, we can extend the definition of \mathbb{P}_σ to any Borel set of runs on cylinders. In particular, the function \mathbb{P}_σ is well-defined for all ω -regular sets of runs, that we will consider in this paper [18]. We extend the Hist notation and let $\text{Hist}(\sigma)$ be the set of histories h such that $\mathbb{P}_\sigma(\text{Cyl}(h)) > 0$. Given a profile σ we denote by $\text{Outcome}(\sigma)$ the set of runs ρ s.t. all prefixes h of ρ belong to $\text{Hist}(\sigma)$. In particular, $\mathbb{P}_\sigma(\text{Outcome}(\sigma)) = 1$. Note that when σ is composed of Dirac strategies then $\text{Outcome}(\sigma)$ is a singleton. The outcome (set of histories) of a strategy $\sigma \in \Gamma_p$, denoted by $\text{Outcome}(\sigma)$ ($\text{Hist}(\sigma)$), is the union of outcomes (set of histories, respectively) of profile σ s.t. $\sigma_p = \sigma$.

Winning conditions To determine the gain of all players in the game \mathcal{G} , we define *winning conditions* that can be interpreted with two kinds of semantics denoted by the symbols \mathbf{S} for the *sure semantics* or \mathbf{A} for the *almost sure semantics*. A winning condition Φ is a subset of $\text{Runs}(\mathcal{G})$ called *winning runs*. From now on, we assume that concurrent games are equipped with a function Φ , called the winning condition, and mapping all players $p \in P$ to a winning condition $\Phi(p)$. A profile σ is \mathbf{A} -winning for $\Phi(p)$ if $\mathbb{P}_\sigma(\Phi) = 1$ which we write $\mathcal{G}, \sigma \models^{\mathbf{A}} \Phi(p)$. A profile σ is \mathbf{S} -winning for $\Phi(p)$ if $\text{Outcome}(\mathcal{G}, \sigma) \subseteq \Phi(p)$ which we write $\mathcal{G}, \sigma \models^{\mathbf{S}} \Phi(p)$. Note that when σ is Dirac, the two semantics coincide: $\mathcal{G}, \sigma \models^{\mathbf{S}} \Phi(p)$ iff $\mathcal{G}, \sigma \models^{\mathbf{A}} \Phi(p)$. The profile σ is \mathbf{A} -winning from h if $h \in \text{Hist}(\mathcal{G}, \sigma)$ and $\mathbb{P}_\sigma(\Phi(p) \mid \text{Cyl}(h)) = \mathbb{P}_\sigma(\Phi(p) \cap \text{Cyl}(h)) / \mathbb{P}_\sigma(\text{Cyl}(h)) = 1$ which we denote $\mathcal{G}, \sigma \models_h^{\mathbf{A}} \Phi(p)$. The profile σ is winning for the sure semantics from h if $\{\rho \in \text{Outcome}(\mathcal{G}, \sigma) \mid h \subseteq_{\text{pref}} \rho\} \subseteq \Phi(p)$, which we denote $\mathcal{G}, \sigma \models_h^{\mathbf{S}} \Phi(p)$. We often omit \mathcal{G} in notations when clear from the context. Most of our definitions and results hold for both semantics and we often state them using the symbol $\star \in \{\mathbf{S}, \mathbf{A}\}$ as in the following definition. Given a semantics $\star \in \{\mathbf{S}, \mathbf{A}\}$, a strategy σ for player p (from a history h) is called \star -winning for player p if for every $\tau \in \Gamma_{-p}$, the profile (σ, τ) is \star -winning for player p (from h). Note that a strategy σ for player p is \mathbf{S} -winning iff $\text{Outcome}(\sigma) \subseteq \Phi(p)$. We often describe winning conditions using standard linear temporal operators \square and \diamond ; e.g. $\square \diamond S$ means the set of runs that visit infinitely often S . See [3] for a formal definition.

A winning condition $\Phi(p)$ is *prefix-independent* if for all $s_1 s_2 \dots \in \Phi(p)$, and all $i \geq 1$: $s_i s_{i+1} \dots \in \Phi(p)$. When $\Phi(p)$ contains all runs that do not visit some designated set $\text{Bad}_p \subseteq S$ of states, we say that $\Phi(p)$ is a *safety condition*. A safety game is a game whose winning condition Φ is such that $\Phi(p)$ is a safety condition for all players p . Without loss of generality, we assume that safety games are so-called *simple* safety games: a safety game $(S, \Sigma, s_{\text{init}}, (\Sigma_p)_{p \in P}, \delta)$ is *simple* iff for all players p , for all $s \in S$: $s \in \text{Bad}_p$ implies that no $s' \notin \text{Bad}_p$ is reachable from s . That is, once the safety condition is violated, then it remains violated forever at all future histories.

► **Example 1.** Let us consider three player-1 strategies in Fig. 2. (i) σ_1 is any strategy that plays a in s_0 ; (ii) σ_2 is any strategy that plays b in s_0 , d in s_1 and f in s_2 ; and (iii) σ_3 is any strategy that plays b in s_0 , d in s_1 , and $0.5f + 0.5g$ in s_2 . Clearly, σ_1 never allows one to reach Trg while some runs respecting σ_2 and σ_3 do (remember that there is no \star -winning strategy in this game). We will see later that the best choice of player 1 (among σ_2, σ_3) depends on the semantics we consider. In the almost-sure semantics, σ_3 is ‘better’ for player 1, because σ_3 is an \mathbf{A} -winning strategy from all histories ending in s_2 , while σ_2 is not. On the other hand, in the sure semantics, playing σ_2 is ‘better’ for player 1 than σ_3 . Indeed, for all player-2 strategies τ , either $\text{Outcome}(\sigma_3, \tau)$ contains only runs that do not reach s_2 , or $\text{Outcome}(\sigma_3, \tau)$ contains at least a run that reaches s_2 , but, in this case, it also contains a run of the form

hs_2^g (because, intuitively, player 1 plays *both* f and g from s_2). So, σ_3 is not winning against any τ , while σ_2 wins at least against a player 2 strategy that plays b' in s_0 , d' in s_1 and f' in s_2 . We formalise these intuitions in the next section.

3 Admissibility

In this section, we define the central notion of the paper: admissibility [4, 8]. Intuitively, a strategy is admissible when it plays ‘as well as possible’. Hence the definition of admissible strategies is based on a notion of domination between strategies: a strategy σ' dominates another strategy σ when σ' wins every time σ does. Obviously, players have no interest in playing dominated strategies, hence admissible strategies are those that are not dominated. Apart from these (classical) definitions, we characterise admissible strategies as those that satisfy two weaker notions: they must be both *strongly cooperative optimal* and play only *locally-admissible moves*. Finally, we discuss important characteristics of admissible strategies that will enable us to perform assume-admissibility synthesis (see Section 4).

In this section, we fix a game \mathcal{G} , a player p , and, following our previous conventions, we denote by Γ_{-p} the set $\{\sigma_{-p} \mid \sigma \in \Gamma\}$.

Admissible strategies We first recall the classical notion of *admissible strategy* [4, 1]. Given two strategies $\sigma, \sigma' \in \Gamma_p$, we say that σ is \star -weakly dominated by σ' , denoted $\sigma \preceq^* \sigma'$, if for all $\tau \in \Gamma_{-p}$: $(\sigma, \tau) \models^* \Phi(p)$ implies $(\sigma', \tau) \models^* \Phi(p)$. This indeed captures the idea that σ' is *not worse* than σ , because it wins (for p) every time σ does. Note that \preceq^* is not anti-symmetric, hence we note $\sigma \approx^* \sigma'$ when σ and σ' are equivalent, i.e. $\sigma \preceq^* \sigma'$ and $\sigma' \preceq^* \sigma$. In other words $\sigma \approx^* \sigma'$ iff for every $\tau \in \Gamma_{-p}$, $(\sigma, \tau) \models^* \Phi(p) \Leftrightarrow (\sigma', \tau) \models^* \Phi(p)$. When $\sigma \preceq^* \sigma'$ but $\sigma' \not\preceq^* \sigma$ we say that σ is \star -dominated by σ' , and we write $\sigma \prec^* \sigma'$. Observe that $\sigma \prec^* \sigma'$ holds if and only if $\sigma \preceq^* \sigma'$ and there exists at least one $\tau \in \Gamma_{-p}$, such that $(\sigma, \tau) \not\models^* \Phi(p)$ and $(\sigma', \tau) \models^* \Phi(p)$. That is, σ' is now *strictly better* than σ . Then, a strategy σ is \star -admissible iff there is no strategy σ' s.t. $\sigma \prec^* \sigma'$, i.e., σ is \star -admissible iff it is not \star -dominated.

► **Example 2.** Let us continue our running example, by formalising the intuitions we have sketched in Example 1. Since σ_1 does not allow to reach the target, while some runs respecting σ_2 and σ_3 do, we have: $\sigma_1 \prec^* \sigma_2$ and $\sigma_1 \prec^* \sigma_3$. Moreover, $\sigma_2 \prec^A \sigma_3$ because σ_3 is **A**-winning from any history that ends in s_2 while σ_2 is not because it does not **A**-win against a player 2 strategy that would always play g' in s_2 (and both strategies behave the same way in s_0 and s_1). On the other hand, $\sigma_3 \prec^S \sigma_2$. In fact, $\text{Outcome}(\sigma_2) \subsetneq \text{Outcome}(\sigma_3)$, so, for all $\tau \in \Gamma_2$, whenever $(\sigma_3, \tau) \models^S \Phi(1)$, all runs of $\text{Outcome}(\sigma_3)$ reach Trg , and we also have $(\sigma_2, \tau) \models^S \Phi(1)$. Moreover, for the strategy $\tau \in \Gamma_2$ that plays f' at s_2 , profile (σ_2, τ) is **S**-winning but not (σ_3, τ) . We will see later that σ_3 is **A**-admissible and σ_2 is **S**-admissible.

Values of histories Before we discuss strongly cooperative optimal and locally admissible strategies, we associate *values* to histories. Let h be a history, and σ be a strategy of player p . Then, the *value of h w.r.t. σ for semantics $\star \in \{\mathbf{S}, \mathbf{A}\}$* is defined as follows. $\chi_\sigma^*(h) = 1$ if σ is \star -winning from h ; $\chi_\sigma^*(h) = 0$ if there are $\tau \in \Gamma_{-p}$ and $\tau' \in \Gamma_{-p}$ s.t. $(\sigma, \tau) \models_h^* \Phi(p)$, and $(\sigma, \tau') \not\models_h^* \Phi(p)$; and -1 otherwise.

Value $\chi_\sigma^*(h) = 1$ corresponds to the case where σ is \star -winning for player p from h (thus, against all possible strategies in Γ_{-p}). When $\chi_\sigma^*(h) = 0$, σ is *not* \star -winning from h (because of τ' in the definition), but the other players can still help p to reach his objective (by playing some τ s.t. $(\sigma, \tau) \models_h^* \Phi(p)$, which exists by definition). Last, $\chi_\sigma^*(h) = -1$ when there is no

hope for p to \star -win, even with the collaboration of the other players. In this case, there is no τ s.t. $(\sigma, \tau) \models_h^* \Phi(p)$. Hence, having $\chi_\sigma^*(h) = -1$ is stronger than saying that σ is not winning—when σ is not winning, we could have $\chi_\sigma^*(h) = 0$ as well.

We define the value of a history h for player p as the best value he can achieve with his different strategies: $\chi_p^*(h) = \max_{\sigma \in \Gamma_p} \chi_\sigma^*(h)$. Last, for $v \in \{-1, 0, 1\}$, let $\text{Val}_{p,v}^*$ be the set of histories h s.t. $\chi_p^*(h) = v$.

Strongly cooperative optimal strategies We are now ready to define *strongly cooperative optimal* (SCO) strategies. Recall that, in the classical setting of turn-based games, admissible strategies are exactly the SCO strategies [8]. We will see that this condition is still necessary but not sufficient in the concurrent setting.

A strategy σ of Player p is \star -SCO at h iff $\chi_\sigma^*(h) = \chi_p^*(h)$; and σ is \star -SCO iff it is \star -SCO at all $h \in \text{Hist}(\sigma)$. Intuitively, when σ is a \star -SCO strategy of Player p , the following should hold: (i) if p has a \star -winning strategy from h (i.e. $\chi_p^*(h) = 1$), then, σ should be \star -winning (i.e. $\chi_\sigma^*(h) = 1$); and (ii) otherwise if p has no \star -winning strategy from h but still has the opportunity to \star -win with the help of other players (hence $\chi_p^*(h) = 0$), then, σ should enable the other players to help p fulfil his objective (i.e. $\chi_\sigma^*(h) = 0$). Observe that when $\chi_p^*(h) = -1$, no continuation of h is \star -winning for p , so $\chi_\sigma^*(h) = -1$ for all strategies σ .

► **Example 3.** Consider again the example in Figure 2. For the almost-sure semantics, we have $\text{Val}_{p,1}^A = \{h \mid \text{last}(h) \in \{s_2, \text{Trg}\}\}$, and $\text{Val}_{p,0}^A = \{h \mid \text{last}(h) \in \{s_0, s_1\}\}$. For the sure semantics, we have: $\text{Val}_{1,1}^S = \{h \mid \text{last}(h) = \text{Trg}\}$, and $\text{Val}_{1,0}^S = \{h \mid \text{last}(h) \neq \text{Trg}\}$. Let us consider again the three strategies σ_1, σ_2 and σ_3 from Example 1. We see that σ_2 is S-SCO but it is not A-SCO because, for all profiles h ending in s_2 : $\chi_{\sigma_2}^A(h) = 0$ while $h \in \text{Val}_{1,1}^A$. On the other hand, σ_3 is A-SCO; but it is not S-SCO. Indeed, one can check that, for all strategies $\tau \in \Gamma_2$: if $\text{Outcome}(\sigma_3, \tau)$ contains a run reaching Trg , then it also contains a run that cycles in s_2 . So, for all such strategies τ , $\text{Outcome}(\sigma_3, \tau) \not\models^S \Phi(1)$, hence $\chi_{\sigma_3}^S(h) = -1$ for all histories that end in s_2 ; while $\chi_p^S(h) = 0$ since $\chi_{\sigma'}^S(h) = 0$ for all Dirac strategies σ' .

Next, let us build a strategy σ'_3 that is A-dominated by σ_3 (hence, not A-admissible), but A-SCO. We let σ'_3 play as σ_3 except that σ'_3 plays c the first time s_1 is visited (hence ensuring that the self-loop on s_1 will be taken after the first visit to s_1). Now, σ_3 is A-dominated by σ'_3 , because (i) σ_3 A-wins every time σ'_3 does; but (ii) σ'_3 does not A-win against the player 2 strategy τ that plays d' only when s_1 is visited for the first time, while σ_3 A-wins against τ . However, σ'_3 is SCO because playing c keeps the value of the history equal to $0 = \chi_1^A(h)$ (intuitively, playing c once does not prevent the other players from helping in the future). As similar example can be built in the S semantics. Thus, **there are \star -SCO strategies which are not admissible**, so, being \star -SCO is not a sufficient criterion for admissibility.

Locally admissible moves and strategies Let us now discuss another criterion for admissibility, which is more *local* in the sense that it is based on a domination between *moves* available to each player after a given history. Let h be a history, and let α and α' be two randomised moves in $\mathcal{D}(\Sigma_p)$. We say that α is *\star -weakly dominated* at h by α' (denoted $\alpha \leq_h^* \alpha'$) iff for all $\sigma \in \Gamma_p$ such that $h \in \text{Hist}(\sigma)$ and $\sigma(h) = \alpha$, there exists $\sigma' \in \Gamma_p$ s.t. $\sigma'(h) = \alpha'$ and $\sigma \preceq^* \sigma'$. Observe that the relation \leq_h^* is not anti-symmetric. We let \simeq_h^* be the equivalence relation s.t. $\alpha \simeq_h^* \beta$ iff $\alpha \leq_h^* \beta$ and $\beta \leq_h^* \alpha$. When $\alpha \leq_h^* \alpha'$ but $\alpha' \not\leq_h^* \alpha$ we say that α is *\star -dominated* at h by α' and denote this by $\alpha <_h^* \alpha'$. When a randomised move α is not \star -dominated at h , we say that α is *\star -admissible* at h . This allows us to define a more local notion of dominated strategy: *a strategy σ of player p is \star -locally-admissible (LA) if $\sigma(h)$ is a \star -admissible move at h , for all histories h .*

► **Example 4.** Consider the Dirac move f and the non-Dirac move $0.5f + 0.5g$ played from s_2 in the example in Figure 2. One can check that $0.5f + 0.5g <_{s_2}^S f$. Indeed, consider a strategy σ s.t. $\sigma(h) = 0.5f + 0.5g$ for some h with $\text{last}(h) = s_2$. Then, playing $\sigma(h)$ from h will never allow Player 1 to reach Trg *surely* at the next step, whatever Player 2 plays; while playing, for instance, f (Dirac move) ensures player 1 to reach Trg surely at the next step, against a Player-2 strategy that plays f' . Thus, σ_2 is S-LA but σ_3 is not.

On the other hand, after every randomised move played in state s_2 , the updated state is s_2 or s_3 from which A-winning strategies exist, thus $f \simeq_h^A g \simeq_h^A \lambda f + (1 - \lambda)g$ for all $\lambda \in [0, 1]$ and all histories h s.t. $\text{last}(h) = s_2$ (so, in particular, $\lambda f + (1 - \lambda)g \leq_h^A f$ and $\lambda f + (1 - \lambda)g \leq_h^A g$). It follows that both σ_2 and σ_3 are A-LA. However, in the long run, player 1 needs to play $\lambda f + (1 - \lambda)g$, with $\lambda \in (0, 1)$, infinitely often in order to A-win. In fact, σ_3 is A-winning from s_2 while σ_2 is not. Thus, **there are \star -LA strategies which are not admissible**, so being \star -LA is not a sufficient criterion for \star -admissibility.

We close this section by several lemmata that allow us to better characterise the notion of LA strategies. First, we observe that, while randomisation might be necessary for winning in certain concurrent games (for example, in Figure 2, no Dirac move allows player 1 to reach Trg *surely* from s_2 , while playing repeatedly f and g with equal probability ensures to reach Trg with probability 1) randomisation is useless when a player wants to play only locally admissible moves. This is shown by the next Lemma (point (i)), saying that, *if a randomised move α plays some action a with some positive probability, then α is dominated by the Dirac move a* . However, this does not immediately allow us to characterise admissible moves: some Dirac moves could be dominated (hence non-admissible), and some non-Dirac moves could be admissible too. Points (ii) and (iii) elucidate this: *among Dirac moves, the non-dominated ones are admissible, and a non-Dirac move is admissible iff all the Dirac moves that occur in its support are admissible and equivalent to each other*.

► **Lemma 5.** *For all histories h and all randomised moves α :*

- (i) *For all $a \in \text{Supp}(\alpha)$: $\alpha \leq_h^* a$;*
- (ii) *Dirac moves that are not \star -dominated at h by another Dirac move are admissible;*
- (iii) *A move α is \star -LA at h iff, for all $a \in \text{Supp}(\alpha)$: (1) a is \star -LA at h ; and (2) $a \simeq_h^* b$ for all $b \in \text{Supp}(\alpha)$.*

► **Example 6.** As we have seen in Example 4, $0.5f + 0.5g <_{s_2}^S f$. Note that a strategy σ' s.t. $\sigma'(h) = 0.5f + 0.5g$ for all h with $\text{last}(h) = s_2$ has value $\chi_{\sigma'}^S(h) = -1$, while $\chi_1^S(h) = 0$.

This example seems to suggest that the local dominance of two moves coincide with the natural order on the values of histories that are obtained when playing those moves (in other words $x <_h^* y$ would hold iff the value of the history obtained by playing x is smaller than or equal to the value obtained by playing y). This is not true for histories of value 0: we have seen that a and b are \leq_h^* -incomparable, yet playing a or b from s_0 yields a history with value 0 in all cases (even when s_1 is reached). The next Lemma gives a precise characterisation of the dominance relation between Dirac moves in terms of values:

► **Lemma 7.** *For all players p , histories h with $\text{last}(h) = s$ and Dirac moves $a, b \in \Sigma_p(s)$: $a \leq_h^* b$ if, and only if the following conditions hold for every $c \in \Sigma_{-p}(s)$ where we write $s_{(a,c)} = \delta(s, (a, c))$ and $s_{(b,c)} = \delta(s, (b, c))$:*

- (i) $\chi_p^*(hs_{(a,c)}) \leq \chi_p^*(hs_{(b,c)})$;
- (ii) if $\chi_p^*(hs_{(a,c)}) = \chi_p^*(hs_{(b,c)}) = 0$ then $s_{(a,c)} = s_{(b,c)}$.

Characterisation and existence of admissible strategies Equipped with our previous results, we can now establish the main results of this section. First, we show that *★-admissible strategies are exactly those that are both ★-LA and ★-SCO* (Theorem 8(i)). Then, we show that *admissible strategies always exist* in concurrent games (Theorem 8(ii)).

► **Theorem 8** (Characterisation and existence of admissible strategies). *The following holds for all strategies σ in a concurrent game with semantics $\star \in \{\mathcal{S}, \mathcal{A}\}$:*

(i) *σ is ★-admissible iff σ is ★-LA and ★-SCO; in the special case of simple safety objectives, if σ is ★-LA then σ is ★-admissible.*

(ii) *there is a ★-admissible strategy σ' such that $\sigma \preceq^{\star} \sigma'$.*

In particular, point (ii) implies that admissible strategies always exist in concurrent games.

► **Example 9.** We consider again the example in Figure 2, and consider strategies σ_2 and σ_3 as defined in Example 1. Remember that these two strategies do their best to reach s_2 , and that, from s_2 , σ_2 plays deterministically f , while σ_3 plays f and g with equal probabilities. From Example 3, we know that σ_2 is \mathcal{S} -SCO but not \mathcal{A} -SCO; while σ_3 is \mathcal{A} -SCO but not \mathcal{S} -SCO. Indeed, we have already argued in Example 2 that σ_2 is not \mathcal{A} -admissible, and that σ_3 is not \mathcal{S} -admissible. However, from Example 4, we know that σ_2 is \mathcal{S} -LA and that σ_3 is \mathcal{A} -LA. So, by Theorem 8, σ_2 is \mathcal{S} -admissible and σ_3 is \mathcal{A} -admissible as expected.

Finally, we close the section by a finer characterisation of ★-admissible strategies. We show that: (i) in the sure semantics, there is always an \mathcal{S} -admissible strategy that plays Dirac moves only; and (ii) in the almost-sure semantics, there is always an \mathcal{A} -admissible strategy that plays Dirac moves only in histories of values 0 or -1 . The difference between the two semantics should not be surprising, as we know already that randomisation is sometimes needed to win (i.e., from histories of value 1) in the almost sure semantics:

► **Proposition 10.** *For all player p strategies σ in a concurrent game with $\star \in \{\mathcal{S}, \mathcal{A}\}$:*

(i) *If σ is \mathcal{A} -admissible then there exists a strategy σ' that plays only Dirac moves in histories of value ≤ 0 such that $\sigma \simeq^{\mathcal{A}} \sigma'$.*

(ii) *If σ is \mathcal{S} -admissible then there exists a Dirac strategy σ' such that $\sigma \simeq^{\mathcal{S}} \sigma'$.*

4 Assume admissible synthesis

In this section we discuss an *assume-admissible synthesis* framework for concurrent games. With classical synthesis, one tries and compute *winning* strategies for all players, i.e., strategies that *always win* against *all possible strategies* of the other players. Unfortunately, it might be the case that such *unconditionally* winning strategies do not exist, as in our example. As explained in the introduction, the assume-admissible synthesis rule relaxes the classical synthesis rule: instead of searching from strategies that win unconditionally, the new rule requires winning against the *admissible strategies* of the other players. So, a strategy may satisfy the new rule while not winning unconditionally. Nevertheless, we claim that winning against admissible strategies is good enough assuming that the players are *rational*; if we assume that players only play strategies that are good for achieving their objectives, they will be playing admissible ones.

The general idea of the assume-admissible synthesis algorithm is to reduce the problem (in a concurrent n -player game) to the synthesis of a winning strategy *in a 2-player zero-sum concurrent game of imperfect information, in the \mathcal{S} -semantics* (even when the original assume-admissible problem is in the \mathcal{A} -semantics), where the objective of player 1 is given by an LTL formula. Such games are solvable using techniques presented in [10].

More precisely, from a concurrent game \mathcal{G} in the semantics $\star \in \{\mathbf{S}, \mathbf{A}\}$ and player p , we build a game \mathcal{G}_p^\star with the above characteristics, which is used to decide the assume-admissible synthesis rule. If such a solution exists, our algorithm constructs a witness strategy. For example, the game \mathcal{G}_1^\star corresponding to the game in Figure 2 is given in Figure 3. The main ingredients for this construction are the following. (i) In \mathcal{G}_p^\star , the protagonist is player p , and the second player is $-p$. (ii) Although randomisation is needed to win in such games in general, we interpret \mathcal{G}_p^\star in the \mathbf{S} -semantics only. In fact, we have seen that for the protagonist, Dirac moves suffice in states of value 0; so the only states where he might need randomisation are those of value 1 (randomisation does not matter if the value is -1 since the objective is lost anyway). Hence we define winning condition to be $\Phi(p) \vee \Diamond \mathbf{Val}_{p,1}^\star$ enabling us to consider only histories of values 0 in \mathcal{G}_p^\star ; and thus hiding the parts of the game where randomisation might be needed. We also prove that we can restrict to Dirac strategies for $-p$ when it comes to admissible strategies. (iii) In order to restrict the strategies to admissible ones, we only allow \star -LA moves in \mathcal{G}_p^\star . These moves can be computed by solving classical 2-player games ([2]) using Lemma 7. For example, in Figure 3, moves c and c' are removed since they are not \mathbf{A} -LA. (iv) Last, since \star -admissible strategies are those that are both \star -LA and \star -SCO (see Theorem 8), we also need to ensure that the players play \star -SCO. This is more involved than \star -LA, as the \star -SCO criterion is not *local*, and requires information about the *sequence of actual moves* that have been played, which cannot be deduced, in a concurrent game, from the sequence of visited states. So, we store, in the states of \mathcal{G}_p^\star , the moves that have been played by all the players to reach the state. For example, in Figure 3, the state labelled by $s_1, (b, b')$ means that \mathcal{G} has reached s_1 , and that the last actions played by the players were b and b' respectively. However, players' strategies must not depend on this extra information since they do not have access to this information in \mathcal{G} either. We thus interpret \mathcal{G}_p^\star as a game of *imperfect information* where all the states labelled by the same state of \mathcal{G} are in the same observation class. Thanks to these constructions, we can finally encode the fact that the players must play \star -SCO strategies in the new objective of the games, which will be given as an LTL formula, as we describe below.

To simplify the presentation, we restrict ourselves to prefix independent winning conditions, although the results can be generalised. Also, to ensure we can effectively solve subproblems mentioned above, we consider ω -regular objectives. The values of the histories depend thus only on their last states, i.e. for all pairs of histories h_1 and h_2 : $\text{last}(h_1) = \text{last}(h_2)$ implies that $\chi_p^\star(h_1) = \chi_p^\star(h_2)$. We thus denote by $\chi_p^\star(s)$ the value $\chi_p^\star(h)$ of all histories h s.t. $\text{last}(h) = s$. Last, we assume that a player cannot play the same action from two different states, i.e. $\forall s_1 \neq s_2, \Sigma_{s_1}(p) \cap \Sigma_{s_2}(p) = \emptyset$. Thus, we say that *a given move a is \star -LA*, meaning that a is \star -LA from all histories ending in the unique state where a is available.

The game \mathcal{G}_p^\star Let us now describe precisely the construction of \mathcal{G}_p^\star . Given an n -player concurrent game $\mathcal{G} = (S, \Sigma, s_{\text{init}}, (\Sigma_p)_{p \in P}, \delta)$ with winning condition Φ considered under semantics $\star \in \{\mathbf{S}, \mathbf{A}\}$, and given a player p , we define the two-player zero-sum concurrent game $\mathcal{G}_p^\star = (\bar{S}, \bar{\Sigma}, \bar{s}_{\text{init}}, (\bar{\Sigma}_p, \bar{\Sigma}_{-p}), \bar{\delta})$ where: (i) $\bar{S} = S \times \bar{\Sigma}^n \cup \{\bar{s}_{\text{init}}\}$; (ii) $\bar{\Sigma}$ is the set of Dirac \star -LA moves in Σ ; (iii) $\bar{s}_{\text{init}} = s_{\text{init}}$ is the initial state; (iv) $\bar{\Sigma}_p$ is such that $\bar{\Sigma}_p(s)$ is the set of Dirac \star -LA moves of p in s , for all $s \in S$; (v) $\bar{\Sigma}_{-p}$ is s.t. for all $s \in S$: $\bar{\Sigma}_{-p}(s)$ is the set of moves \mathbf{a} of $-p$ in s s.t. for all $q \neq p$, a_q is a Dirac \star -LA move; (vi) $\bar{\delta}$ updates the state according to δ , remembering the last actions played: $\bar{\delta}(\bar{s}_{\text{init}}, \mathbf{b}) = (\delta(s_{\text{init}}, \mathbf{b}), \mathbf{b})$ and $\bar{\delta}((s, \mathbf{a}), \mathbf{b}) = (\delta(s, \mathbf{b}), \mathbf{b})$ for all $s \in S$. Note that the game \mathcal{G}_p^\star depends on whether $\star = \mathbf{A}$ or $\star = \mathbf{S}$ because the two semantics yield different sets of LA-moves. However, we interpret \mathcal{G}_p^\star in the sure semantics, so both players can play Dirac strategies only in \mathcal{G}_p^\star .

Let us now explain how we obtain an imperfect information game by defining an *observation function* \mathfrak{o} . Note that histories in \mathcal{G}_p^* are of the form: $\bar{h} = \overline{s_{\text{init}}}(s_1, \mathbf{a}_1)(s_2, \mathbf{a}_2) \cdots (s_n, \mathbf{a}_n)$. Then, let $\mathfrak{o} : \bar{S} \rightarrow S$ be the mapping that, intuitively, projects moves away from states. For example, in Figure 3, states with observation s_0 are in the dashed rectangle. That is: $\mathfrak{o}(s, \mathbf{a}) = s$ for all states s , and $\mathfrak{o}(\overline{s_{\text{init}}}) = s_{\text{init}}$. We extend \mathfrak{o} to histories recursively: $\mathfrak{o}(\overline{s_{\text{init}}}) = s_{\text{init}}$ and $\mathfrak{o}(h(s_n, \mathbf{a}_n)) = \mathfrak{o}(h)s_n$. To make \mathcal{G}_p^* a game of imperfect information, we request that, in \mathcal{G}_p^* , players play only strategies σ s.t. $\sigma(h_1) = \sigma(h_2)$ whenever $\mathfrak{o}(h_1) = \mathfrak{o}(h_2)$.

We relate the strategies in the original game \mathcal{G} with the strategies in \mathcal{G}_p^* , which we need to extract admissible strategies in \mathcal{G} from the winning strategies in \mathcal{G}_p^* and thus perform assume-admissible synthesis. First, given a player p strategy σ in \mathcal{G} (i.e., $\sigma \in \Gamma_p(\mathcal{G})$), we say that a strategy $\bar{\sigma} \in \Gamma_p^{\text{det}}(\mathcal{G}_p^*)$ is a *realisation* of σ iff: (i) $\bar{\sigma}$ is Dirac; and (ii) $\bar{\sigma}(h) \in \text{Supp}(\sigma(h))$ for all h . Note that every \star -LA strategy $\sigma \in \Gamma_i(\mathcal{G})$ admits realisations $\bar{\sigma}$ in $\Gamma_i(\mathcal{G}_p^*)$. Second, given a player p Dirac strategy σ in \mathcal{G}_p^* (i.e., $\sigma \in \Gamma_p^{\text{det}}(\mathcal{G}_p^*)$) we say that $\hat{\sigma} \in \Gamma_p(\mathcal{G})$ is an *extension* of σ iff, for all $h \in \text{Hist}(\mathcal{G}_p^*, \sigma)$: $\hat{\sigma}(\mathfrak{o}(h)) = \sigma(h)$.

The assume-admissible synthesis technique As explained above, the assume-admissible rule boils down to computing a winning strategy $\bar{\sigma}$ for player p in \mathcal{G}_p^* w.r.t. to winning condition $\Phi_{\mathcal{G}_p^*}$, and extracting, from $\bar{\sigma}$, the required admissible strategy in \mathcal{G} .

We will now formally define $\Phi_{\mathcal{G}_p^*}$. Let p be a player (in \mathcal{G}); and let us denote by $\text{st}(a)$ the (unique) state from which a is available, for all actions a . We define AfterHelpMove_p^* as

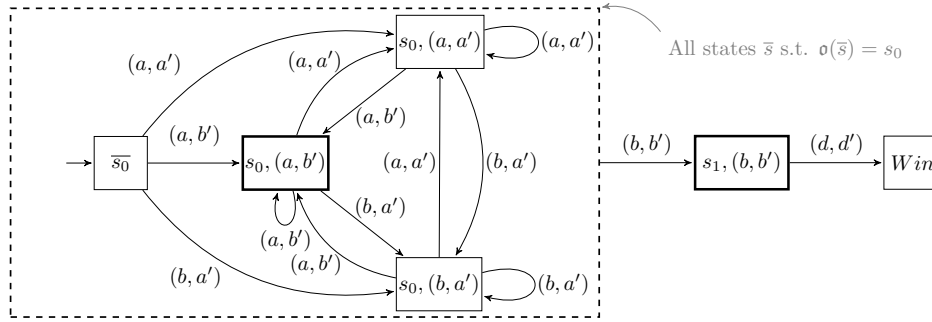
$$\text{AfterHelpMove}_p^* = \{(s, \mathbf{a}) \in \bar{S} \mid \exists s' \in \text{Succ}(\text{st}(a_p), a_p) : \chi_p^*(s') \geq 0 \wedge s' \neq s \wedge \chi_p^*(s) = 0\}.$$

That is, when $(s, \mathbf{a}) \in \text{AfterHelpMove}_p^*$, in \mathcal{G} , player p has played a_p from $\text{st}(a_p)$ and, due to player $-p$'s choice, \mathcal{G} has reached s . However, with another choice of player $-p$, the game could have moved to a different state s' from which $-p$ can help p to win as $\chi_p^*(s') \geq 0$. Intuitively, in runs that visit states of value 0 infinitely often, states from AfterHelpMove_p^* should be visited infinitely often for player i to play SCO, i.e. such runs might not be winning, but this cannot be blamed on player p who has sought repeatedly the collaboration of the other players to enforce his objective. Observe further that the definition of this predicate requires the labelling of the states (by actions) we have introduced in \mathcal{G}_p^* . For example, in Figure 3, $\text{AfterHelpMove}_2^* = \{(s_0, (a, b')), (s_1, (b, b'))\}$. We let $\Phi_0^*(p) = \diamond \neg \text{Val}_{p,0}^* \vee \Phi(p) \vee \square \diamond \text{AfterHelpMove}_p^*$ and $\Phi_1^*(p) = (\diamond \text{Val}_{p,1}^*) \rightarrow \Phi(p)$. Let us define $\Phi_{\mathcal{G}_p^*} = \left(\bigwedge_{q \neq p} \Phi_0^*(q) \wedge \Phi_1^*(q) \right) \rightarrow (\Phi(p) \vee \diamond \text{Val}_{p,1}^*)$.

► **Theorem 11** (Assume-admissible synthesis). *Player p has a \star -admissible strategy σ that is \star -winning against all player $-p$ \star -admissible strategies in \mathcal{G} iff Player p has an \mathcal{S} -winning strategy in \mathcal{G}_p^* for the objective $\Phi_{\mathcal{G}_p^*}$. Such a \star -admissible strategy σ can be effectively computed (from any player p \mathcal{S} -winning strategy in \mathcal{G}_p^*).*

Let us explain how we build a strategy in \mathcal{G} with the desired properties, from any player p strategy enforcing $\Phi_{\mathcal{G}_p^*}$ in \mathcal{G}_p^* . Remember that \mathcal{G}_p^* ensures that the players play \star -LA moves only. We will use $\Phi_{\mathcal{G}_p^*}$ to make sure that, when SCO strategies are played by $-p$ (relying on the extra information we have encoded in the states), then p reaches a state of value 1. First, consider $\Phi_0^*(q)$ for $q \neq p$. Runs that satisfy this formula are either those that visit states of value 0 only finitely often ($\diamond \neg \text{Val}_{q,0}^*$); or those that stay in states of value 0, in which case they must be either winning ($\Phi(q)$) or visit infinitely often states where Player q could have been helped by the other players ($\square \diamond \text{AfterHelpMove}_q^*$). This is a necessary condition on runs visiting only value 0 states for the strategy to be SCO. Next, observe that

$\Phi_1^*(q)$ states that *if* a history of value 1 is entered *then* Player q must win. This allows us to understand the left part of the implication in $\Phi_{\mathcal{G}_p^*}$: the implication can be read as ‘if all other players play a \star -admissible strategy, then either p should win ($\Phi(p)$) or a state of value 1 for player p should eventually be visited ($\Diamond \text{Val}_{p,1}^*$)’. Then a strategy $\hat{\sigma}$ (in \mathcal{G}) that wins against admissible strategies can be extracted from a winning strategy $\bar{\sigma}$ (in \mathcal{G}_p^*) in a straightforward way, *except* when $\bar{\sigma}$ enforces to reach a state of value 1 ($\Diamond \text{Val}_{p,1}^*$ in $\Phi_{\mathcal{G}_p^*}$). In this case, σ cannot follow $\bar{\sigma}$, but must rather switch to a *winning* strategy, which: (i) is guaranteed to exist since the state that has been reached has value 1; and (ii) can be computed using classical techniques [10]. The strategy $\hat{\sigma}$ is not necessarily admissible but by Theorem 8 (i), there is an admissible strategy σ with $\hat{\sigma} \preceq^* \sigma$. By weak domination, σ wins against more profiles than $\hat{\sigma}$, in particular, it wins against the profiles of admissible strategies of the other players.



■ **Figure 3** The game \mathcal{G}_1^A obtained from the game in Figure 2. Bold states $(s_0, (a, b'))$ and $(s_1, (b, b'))$ are the states of AfterHelpMove_2^A . There is a (b, b') -labelled transition from all states in the dashed rectangle to $(s_1, (b, b'))$.

► **Example 12.** In our running example, observe that $-\text{Val}_{2,0}^A = \text{Val}_{2,1}^A = \{Win\}$ since there is no state of value -1 in \mathcal{G} . Hence, $\Phi(2) = \Diamond Win = \Diamond \text{Val}_{2,1}^A = \Diamond -\text{Val}_{2,0}^A$. Finally, $\text{AfterHelpMove}_2^A = \{(s_0, (a, b')), (s_1, (b, b'))\}$, so, after simplification: $\Phi_{\mathcal{G}_1^A} = [\Diamond Win \vee \square \Diamond ((s_0, (a, b')) \vee (s_1, (b, b')))] \rightarrow \Diamond Win$. Thus, to win in \mathcal{G}_1^A (under the *sure* semantics), player 1 must ensure to reach *Win* as long as player 2 visits the set of *bold* states in Figure 3 infinitely often. A winning strategy $\bar{\sigma}$ in \mathcal{G}_1^A consists in (eventually) always playing b from all states in the dashed rectangle; and d from $(s_1, (b, b'))$. Observe that this strategy is compatible with \mathbf{o} . From $\bar{\sigma}$, we can extract an admissible player 1 strategy in \mathcal{G} : always play b in s_0 ; always play d in s_1 ; and play a winning strategy from s_2 (which is of value 1), for instance: always play $0.5f + 0.5g$ from s_2 like σ_3 does.

We conclude by a remark on games with simple safety objectives.

► **Remark.** In the case of simple safety games, the situation is much simpler. We have seen in Theorem 8 that, for simple safety objectives, \star -LA strategies are exactly the admissible strategies. So, can simply build \mathcal{G}_p from \mathcal{G} by pruning the actions which are not \star -LA (the labelling by actions is not necessary anymore since its sole purpose is to enforce SCO), and look for a player p winning strategy in the resulting game.

References

- 1 Brandenburger Adam, Friedenberg Amanda, H Jerome, et al. Admissibility in games. *Econometrica*, 2008.

- 2 Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002. URL: <http://doi.acm.org/10.1145/585265.585270>, doi:10.1145/585265.585270.
- 3 Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking (Representation and Mind Series)*. The MIT Press, 2008.
- 4 Dietmar Berwanger. Admissibility in infinite games. In *STACS 2007, 24th Annual Symposium on Theoretical Aspects of Computer Science, Aachen, Germany, February 22-24, 2007, Proceedings*, number 4393 in Lecture Notes in Computer Science, pages 188–199. Springer, 2007. URL: http://dx.doi.org/10.1007/978-3-540-70918-3_17, doi:10.1007/978-3-540-70918-3_17.
- 5 Patricia Bouyer, Romain Brenguier, Nicolas Markey, and Michael Ummels. Nash equilibria in concurrent games with Büchi objectives. In *Proceedings of the 31st Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'11)*, volume 13 of *Leibniz International Proceedings in Informatics*, pages 375–386, Mumbai, India, dec 2011. Leibniz-Zentrum für Informatik.
- 6 Romain Brenguier, Guillermo A. Pérez, Jean-François Raskin, and Ocan Sankur. Admissibility in Quantitative Graph Games. In Akash Lal, S. Akshay, Saket Saurabh, and Sandeep Sen, editors, *36th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2016)*, volume 65 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 42:1–42:14, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: <http://drops.dagstuhl.de/opus/volltexte/2016/6877>, doi:<http://dx.doi.org/10.4230/LIPIcs.FSTTCS.2016.42>.
- 7 Romain Brenguier, Jean-François Raskin, and Ocan Sankur. Assume-admissible synthesis. In Luca Aceto and David de Frutos-Escrig, editors, *CONCUR*, volume 42 of *LIPIcs*, pages 100–113. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015. doi:10.4230/LIPIcs.CONCUR.2015.100.
- 8 Romain Brenguier, Jean-François Raskin, and Ocan Sankur. Assume-admissible synthesis. *Acta Inf.*, 54(1):41–83, 2017. URL: <http://dx.doi.org/10.1007/s00236-016-0273-2>, doi:10.1007/s00236-016-0273-2.
- 9 Romain Brenguier, Jean-François Raskin, and Mathieu Sassolas. The complexity of admissibility in omega-regular games. In *CSL-LICS '14, 2014*. ACM, 2014. URL: <http://doi.acm.org/10.1145/2603088.2603143>.
- 10 Krishnendu Chatterjee, Luca de Alfaro, and Thomas A. Henzinger. Qualitative concurrent parity games. *ACM Trans. Comput. Log.*, 12(4):28:1–28:51, 2011. URL: <http://doi.acm.org/10.1145/1970398.1970404>, doi:10.1145/1970398.1970404.
- 11 Krishnendu Chatterjee and Thomas A. Henzinger. Assume-guarantee synthesis. In *Tools and Algorithms for the Construction and Analysis of Systems, 13th International Conference, TACAS 2007, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2007 Braga, Portugal, March 24 - April 1, 2007, Proceedings*, volume 4424 of *Lecture Notes in Computer Science*, pages 261–275. Springer, 2007.
- 12 Werner Damm and Bernd Finkbeiner. Automatic compositional synthesis of distributed systems. In *FM 2014: Formal Methods - 19th International Symposium, Singapore, May 12-16, 2014. Proceedings*, volume 8442 of *Lecture Notes in Computer Science*, pages 179–193. Springer, 2014.
- 13 Luca de Alfaro, Thomas A. Henzinger, and Orna Kupferman. Concurrent reachability games. *Theor. Comput. Sci.*, 386(3):188–217, 2007. URL: <http://dx.doi.org/10.1016/j.tcs.2007.07.008>, doi:10.1016/j.tcs.2007.07.008.
- 14 Marco Faella. Admissible strategies in infinite games over graphs. In *MFCS 2009*, volume 5734 of *Lecture Notes in Computer Science*, pages 307–318. Springer, 2009.

- 15 Dana Fisman, Orna Kupferman, and Yoav Lustig. Rational synthesis. In *Tools and Algorithms for the Construction and Analysis of Systems, 16th International Conference, TACAS 2010, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2010, Paphos, Cyprus, March 20-28, 2010. Proceedings*, volume 6015 of *Lecture Notes in Computer Science*, pages 190–204. Springer, 2010.
- 16 Orna Kupferman, Giuseppe Perelli, and Moshe Y. Vardi. Synthesis with rational environments. In *Multi-Agent Systems - 12th European Conference, EUMAS 2014, Prague, Czech Republic, December 18-19, 2014, Revised Selected Papers*, pages 219–235. Springer, 2014.
- 17 Jean-François Raskin, Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. Algorithms for omega-regular games with imperfect information. *Logical Methods in Computer Science*, 3(3), 2007. URL: [http://dx.doi.org/10.2168/LMCS-3\(3:4\)2007](http://dx.doi.org/10.2168/LMCS-3(3:4)2007), doi:10.2168/LMCS-3(3:4)2007.
- 18 Moshe Y Vardi. Automatic verification of probabilistic concurrent finite state programs. In *Foundations of Computer Science, 1985., 26th Annual Symposium on*, pages 327–338. IEEE, 1985.

A Proofs for Section 2

We note the following property of winning strategies.

► **Lemma 13.** *Given $\star \in \{S, A\}$, $\sigma \models^{\star} \Phi(p)$ iff $\sigma \models_h^{\star} \Phi(p)$ for every $h \in \text{Hist}(\sigma)$.*

A.1 Proof of Lemma 13

Proof. The only non-trivial implication is: if $\sigma \models^A \Phi(p)$ then $\sigma \models_h^A \Phi(p)$ for every $h \in \text{Hist}(\sigma)$ which we show by contraposition. Assume there exists $h \in \text{Hist}(\sigma)$ such that $\sigma \not\models_h^A \Phi(p)$. This means that $\mathbb{P}_{\sigma}(\text{Cyl}(h)) > 0$ and $\mathbb{P}_{\sigma}(\Phi(p) \cap \text{Cyl}(h)) / \mathbb{P}_{\sigma}(\text{Cyl}(h)) < 1$. It follows that $\sigma \not\models_h^A \Phi(p)$. ◀

B Proofs for Section 3

B.1 Proof of Lemma 5

Proof. Proof of (i): Take a strategy $\sigma \in \Gamma_p$ such that $h \in \text{Hist}(\sigma)$ and $\sigma(h) = \alpha$. Define σ' the strategy that plays as σ except in h where it plays a instead of α . One show that $\sigma \preceq^{\star} \sigma'$. Consider any $\tau \in \Gamma_{-p}$ such that $(\sigma, \tau) \models_h^A \Phi(p)$. This means that $\mathbb{P}_{(\sigma, \tau)}(\Phi(p) \mid \text{Cyl}(h)) = 1$, so, in particular $\mathbb{P}_{(\sigma, \tau)}(\Phi(p) \mid \text{Cyl}(h\delta(s, a, \tau(h)))) = 1$, for all $a \in \text{Supp}(\alpha)$. Since σ and σ' are identical in all other histories, including those extending $h\delta(s, a, \tau(h))$, we deduce that $(\sigma', \tau) \models_h^A \Phi(p)$. The proof for the sure semantics is similar.

Proof of (ii): If a Dirac move a is \star -dominated at h by a move α' then by (i) it is \star -dominated at h by a Dirac move $a' \in \text{Supp}(\alpha')$. This shows (ii) by contraposition.

Proof of (iii): Assume α is a \star -LA move at h . By (i), for every $a \in \text{Supp}(\alpha)$, $\alpha \leq_h^{\star} a$ and $\alpha \not\prec_h^{\star} a$ (because α is a \star -LA move at h) so $\alpha \simeq_h^{\star} a$. Hence all the elements of $\text{Supp}(\alpha)$ are equivalent and \star -LA at h .

Assume now that $\forall a \in \text{Supp}(\alpha)$, a is \star -LA at h and $\forall b \in \text{Supp}(\alpha)$, $a \simeq_h^{\star} b$. We take $a \in \text{Supp}(\alpha)$. By Lemma 5 it holds that $\alpha \leq_{\sigma}^{\star} a$; so it remains to show that $a \leq_h^{\star} \alpha$. Let σ be such that $\sigma(h) = a$. We construct a strategy σ' such that $\sigma'(h) = \alpha$ and $\sigma \preceq^{\star} \sigma'$. For every $b \simeq_h^{\star} a$, there exists a strategy σ_b such that $\sigma_b(h) = b$ and $\sigma \preceq^{\star} \sigma_b$. We construct σ' as follows. We start with $\sigma' = \sigma$, we set $\sigma'(h) = \alpha$, for every $b \in \text{Supp}(\alpha)$ and for every $s' \in \text{Succ}(s, b)$, we do $\sigma' \leftarrow \sigma' \langle hs' \leftarrow \sigma_b \rangle$. This shows that $a \leq_h^{\star} \alpha$. We conclude that α is \star -LA at h . ◀

B.2 Proof of Lemma 7

Proof. Direction $a \leq_h^{\star} b \Rightarrow (\text{i}) \wedge (\text{ii})$: Assume that $a \leq_h^{\star} b$. Let $\sigma \in \Gamma_p$ such that $h \in \text{Hist}(\sigma)$ and $\sigma(h) = a$ and strategy $\sigma' \in \Gamma_p$ such that $\sigma'(h) = b$ and $\sigma \preceq^{\star} \sigma'$.

Proof of (i): We take $c \in \Sigma_p(s)$, and show that $\chi_p^{\star}(hs_{(a,c)}) \leq \chi_p^{\star}(hs_{(b,c)})$. For every $\tau \in \Gamma_{-p}$, if $(\sigma, \tau) \models^{\star} \Phi(p)$ then $(\sigma', \tau) \models^{\star} \Phi(p)$. In particular, if $(\sigma, \tau) \models_{hs_{(a,c)}}^{\star} \Phi(p)$ then $(\sigma', \tau) \models_{hs_{(b,c)}}^{\star} \Phi(p)$. We deduce that σ' is winning from $hs_{(b,c)}$ if σ is winning from $hs_{(a,c)}$. As σ is arbitrary this shows that $\chi_p^{\star}(hs_{(a,c)}) = 1 \Rightarrow \chi_p^{\star}(hs_{(b,c)}) = 1$. Dually if $\chi_p^{\star}(hs_{(b,c)}) = -1$ then σ' is losing from $hs_{(b,c)}$ and so is σ from $hs_{(a,c)}$. This shows the implication $\chi_p^{\star}(hs_{(b,c)}) = -1 \Rightarrow \chi_p^{\star}(hs_{(a,c)}) = -1$. These two implications yield $\chi_p^{\star}(hs_{(a,c)}) \leq \chi_p^{\star}(hs_{(b,c)})$.

Proof of (ii): We show the contrapositive, consider c such that $s_{(a,c)} \neq s_{(b,c)}$ and show that $\chi_p^{\star}(hs_{(a,c)})$ and $\chi_p^{\star}(hs_{(b,c)})$ cannot be both equal to 0. For this purpose we assume that $\chi_p^{\star}(hs_{(a,c)}) = 0$ and show that $\chi_p^{\star}(hs_{(b,c)}) = 1$. Since σ is chosen arbitrary we take it such that $\chi_{\sigma}^{\star}(hs_{(a,c)}) = \chi_p^{\star}(hs_{(a,c)}) = 0$. In particular there exists $\tau \in \Gamma_{-p}$ such that

$(\sigma, \tau) \models_{hs(a,c)}^* \Phi(p)$. Let τ' be an arbitrary player $-p$ strategy. By $\sigma \preceq^* \sigma'$, we have $(\sigma, \tau') \models^* \Phi(p)$ implies $(\sigma', \tau') \models^* \Phi(p)$. Since $\tau'(hs(a,c) \leftarrow \tau)$ and τ' are equal on histories incomparable to $hs(a,c)$, it then holds that $(\sigma', \tau') \models_{hs(b,c)}^* \Phi(p)$ for every arbitrary profile τ' . We have shown that σ' is winning from $hs(b,c)$ and hence that $\chi_p^*(hs(b,c)) = 1$ as claimed.

Direction (i) \wedge (ii) $\Rightarrow a \preceq_h^* b$: We assume that (i) and (ii) hold and show that $a \preceq_h^* b$. Let σ be such that $h \in \text{Hist}(\sigma)$ and $\sigma(h) = a$. We define σ' such that $\sigma'(h) = b$ and $\sigma \preceq^* \sigma'$ in several steps. We start with $\sigma' = \sigma$. We set $\sigma'(h) = b$. For every $s' \in \text{Succ}(s, b) \setminus \text{Succ}(s, a)$ we take a strategy $\sigma_{s'}$ such that $\chi_{\sigma_{s'}}^*(hs') = \chi_p^*(hs')$ and update σ' to $\sigma' \langle hs' \leftarrow \sigma_{s'} \rangle$. For every $s' \in \text{Succ}(s, b) \cap \text{Succ}(s, a)$ such that $\chi_{\sigma'}^*(hs') < \chi_p^*(hs')$ we do the same operation. Note that for all states $s' \in \text{Succ}(s, b) \cap \text{Succ}(s, a)$ such that $\chi_{\sigma'}^*(hs') = \chi_p^*(hs')$ the strategies σ' and σ coincide from hs' . Now we show that $\sigma \preceq^* \sigma'$. We take an arbitrary $\tau \in \Gamma_{-p}$ such that $(\sigma, \tau) \models^* \Phi(p)$ and show that $(\sigma', \tau) \models^* \Phi(p)$. Assume that $h \in \text{Hist}(\sigma, \tau)$, otherwise (σ, τ) and (σ', τ) behave exactly the same. Let $\gamma = \tau(h)$. We show that $(\sigma', \tau) \models_h^* \Phi(p)$. This is equivalent to show that $(\sigma', \tau) \models_{hs(b,c)}^* \Phi(p)$ for every $c \in \text{Supp}(\gamma)$. Let $c \in \text{Supp}(\gamma)$. If $s(b,c) = s(a,c)$ and $\chi_{\sigma'}^*(hs(b,c)) = \chi_p^*(hs(b,c))$ then σ and σ' play the same from $hs(a,c) = hs(b,c)$ and so $(\sigma', \tau) \models_{hs(b,c)}^* \Phi(p)$. If $s(b,c) \neq s(a,c)$ then by (ii) and (i), either $\chi_p^*(hs(b,c)) = \chi_p^*(hs(a,c)) = -1$, or $\chi_p^*(hs(b,c)) = \chi_p^*(hs(a,c)) = 1$, or $\chi_p^*(hs(b,c)) > \chi_p^*(hs(a,c))$. The first case is not possible since τ witnesses that $\chi_p^*(hs(a,c)) > 0$. In the second case, by construction, we have $\chi_{\sigma'}^*(hs(b,c)) = 1$, thus $(\sigma', \tau) \models_{hs(b,c)}^* \Phi(p)$. Assume the third case, which means $\chi_p^*(hs(b,c)) > \chi_p^*(hs(a,c))$. Then $\chi_{\sigma'}^*(hs(b,c)) = \chi_p^*(hs(b,c)) > \chi_p^*(hs(a,c)) \geq 0$ (the first equality is due to the construction of σ' , the last inequality is due to the fact that $(\sigma, \tau) \models_{hs(a,c)}^* \Phi(p)$). We deduce that σ' is winning from $hs(b,c)$ and hence that $(\sigma', \tau) \models_{hs(b,c)}^* \Phi(p)$. This concludes the proof of $(\sigma', \tau) \models_h^* \Phi(p)$. As σ' agrees with σ on histories for which h is not a prefix, this implies that $(\sigma', \tau) \models^* \Phi(p)$. \blacktriangleleft

B.3 Lemma 14

We turn our attention to randomised moves in general, and show, again, how the notion of \star -LA equivalence between \star -LA moves relates to values of histories. Let h be a history ending in s , and assume that player 1 has the choice between two randomised moves α and β which are \star -LA and equivalent ($\alpha \simeq_h^* \beta$). Then, the probability of ending up in some state s' will be the same by playing α or β , provided hs' has value 0 (for player p). A similar property exists for histories of value 1: the probability of producing a history of value 1 is the same under α and β :

► **Lemma 14.** *Let h be a history s.t. $\text{last}(h) = s$, and let α and β be two randomised moves in $\mathcal{D}(\Sigma_p(s))$ s.t. $\alpha \simeq_h^* \beta$. Then, for all $\gamma \in \mathcal{D}(\Sigma_{-p}(s))$:*

- (i) $\delta_r(s, (\alpha, \gamma))(s') = \delta_r(s, (\beta, \gamma))(s')$ for all s' s.t. $hs' \in \text{Val}_{p,0}^*$; and
- (ii) $\delta_r(s, (\alpha, \gamma))(S_1) = \delta_r(s, (\beta, \gamma))(S_1)$, with $S_1 = \{s' \mid hs' \in \text{Val}_{p,1}^*\}$.

Proof. We first prove the properties for Dirac moves $\alpha = a$ and $\beta = b$ such that $a \simeq_h^* b$. By Lemma 7, $\chi_{h\delta(s,a,c)}^p = \chi_{h\delta(s,b,c)}^p$ for every $c \in \Sigma_{-p}(s)$ and if this value is 0 it further holds that: $\delta(s, a, c) = \delta(s, b, c)$. Then for every s' such that $\chi_{hs'}^p = 0$, for every $c \in \text{Supp}(\gamma)$, it holds that $\delta_r(s, (a, c))(s') = \delta_r(s, (b, c))(s')$. Thus (i) holds when $\alpha = a$ and $\beta = b$ are Dirac

moves:

$$\begin{aligned}\delta_r(s, (a, \gamma))(s') &= \sum_{c \in \text{Supp}(\gamma)} \gamma(c) \delta_r(s, (a, c))(s') \\ &= \sum_{c \in \text{Supp}(\gamma)} \gamma(c) \delta_r(s, (b, c))(s') \\ &= \delta_r(s, (b, \gamma))(s').\end{aligned}$$

Next, we know that, for all Dirac moves $c \in \Sigma_{-p}(s)$, $\delta_r(s, (a, c))(\delta(s, a, c)) = \delta_r(s, (b, c))(\delta(s, b, c)) = 1$, by definition of Dirac moves. We use this to establish point (ii), again when $\alpha = a$ and $\beta = b$ are Dirac moves:

$$\begin{aligned}\delta_r(s, (a, \gamma))(S_1) &= \sum_{s' | h s' \in \text{Val}_{p,1}^*} \delta_r(s, (a, \gamma))(s') \\ &= \sum_{c | h \delta(s, a, c) \in \text{Val}_{p,1}^*} \gamma(c) \delta_r(s, (a, c))(\delta(s, a, c)) \\ &= \sum_{c | h \delta(s, b, c) \in \text{Val}_{p,1}^*} \gamma(c) \delta_r(s, (b, c))(\delta(s, b, c)) \\ &= \sum_{s' | h s' \in \text{Val}_{p,1}^*} \delta_r(s, (b, \gamma))(s') \\ &= \delta_r(s, (b, \gamma))(S_1).\end{aligned}$$

Notice that the third inequality follows from $a \simeq_h^* b$ and Lemma 7, since we have $\{c \mid h \delta(s, a, c) \in \text{Val}_{p,1}^*\} = \{c \mid h \delta(s, b, c) \in \text{Val}_{p,1}^*\}$.

To finish the proof, we establish both properties for general randomised moves α and β . We know that α 's and β 's supports are equivalent Dirac moves by Lemma 5. We start by selecting some $b \in \text{Supp}(\beta)$, in order to reuse the results above. We have:

$$\begin{aligned}\delta_r(s, (\alpha, \gamma))(s') &= \sum_{a \in \text{Supp}(\alpha)} \delta_r(s, (a, \gamma))(s') \alpha(a) \\ &= \sum_{a \in \text{Supp}(\alpha)} \delta_r(s, (b, \gamma))(s') \alpha(a) && \text{as proved above} \\ &= \delta_r(s, (b, \gamma))(s') \sum_{a \in \text{Supp}(\alpha)} \alpha(a) \\ &= \delta_r(s, (b, \gamma))(s').\end{aligned}$$

So, summing up:

$$\delta_r(s, (\alpha, \gamma))(s') = \delta_r(s, (b, \gamma))(s'). \quad (1)$$

However:

$$\begin{aligned}\delta_r(s, (\beta, \gamma))(s') &= \sum_{b \in \text{Supp}(\beta)} \delta_r(s, (b, \gamma))(s') \beta(b) \\ &= \sum_{b \in \text{Supp}(\beta)} \delta_r(s, (\alpha, \gamma))(s') \beta(b) && \text{by (1)} \\ &= \delta_r(s, (\alpha, \gamma))(s') \sum_{b \in \text{Supp}(\beta)} \beta(b) \\ &= \delta_r(s, (\alpha, \gamma))(s').\end{aligned}$$

Therefore establishing (i).

Finally, we proceed similarly for S_1 . Again, we start by fixing $b \in \text{Supp}(\beta)$. We have:

$$\begin{aligned}
 \delta_r(s, (\alpha, \gamma))(S_1) &= \sum_{s' | hs' \in \text{Val}_{p,1}^*} \delta_r(s, (\alpha, \gamma))(s') \\
 &= \sum_{a \in \text{Supp}(\alpha)} \left(\alpha(a) \sum_{s' | hs' \in \text{Val}_{p,1}^*} \delta_r(s, (a, \gamma))(s') \right) \\
 &= \sum_{a \in \text{Supp}(\alpha)} \left(\alpha(a) \sum_{s' | hs' \in \text{Val}_{p,1}^*} \delta_r(s, (b, \gamma))(s') \right) \quad \text{as proved above} \\
 &= \left(\sum_{s' | hs' \in \text{Val}_{p,1}^*} \delta_r(s, (b, \gamma))(s') \right) \left(\sum_{a \in \text{Supp}(\alpha)} \alpha(a) \right) \\
 &= \sum_{s' | hs' \in \text{Val}_{p,1}^*} \delta_r(s, (b, \gamma))(s').
 \end{aligned}$$

So, summing up:

$$\delta_r(s, (\alpha, \gamma))(S_1) = \sum_{s' | hs' \in \text{Val}_{p,1}^*} \delta_r(s, (b, \gamma))(s'). \quad (2)$$

However,

$$\begin{aligned}
 \delta_r(s, (\beta, \gamma))(S_1) &= \sum_{s' | hs' \in \text{Val}_{p,1}^*} \delta_r(s, (\beta, \gamma))(s') \\
 &= \sum_{b \in \text{Supp}(\beta)} \left(\beta(b) \sum_{s' | hs' \in \text{Val}_{p,1}^*} \delta_r(s, (b, \gamma))(s') \right) \\
 &= \sum_{b \in \text{Supp}(\beta)} \beta(b) \delta_r(s, (\alpha, \gamma))(S_1) \quad \text{by (2)} \\
 &= \delta_r(s, (\alpha, \gamma))(S_1) \sum_{b \in \text{Supp}(\beta)} \beta(b) \\
 &= \delta_r(s, (\alpha, \gamma))(S_1)
 \end{aligned}$$

Therefore establishing (ii). ◀

B.4 Proof of Theorem 8

To establish Theorem 8, we need an ancillary lemma, which relates the notions of \star -admissibility and the notions of \star -LA. To this end, we draw a link between (global) equivalence of strategies (in terms of \approx^*) and the local equivalence of moves (in terms of \simeq_h^*). For all player p strategies σ and all $v \in \{-1, 0, 1\}$, let us define $\text{Hist}_v(\sigma)$ as $\text{Hist}(\sigma) \cap \text{Val}_{p,v}^*$, i.e., the set of histories of σ that have value v . Then:

► **Lemma 15.** *Let σ and σ' be two player p strategies s.t.: (i) σ is \star -LA; and (ii) for all histories h : $\chi_p^*(h) = 1$ implies that σ and σ' are \star -winning from h . Then, the following conditions are equivalent:*

(1) $\sigma(h) \simeq_h^* \sigma'(h)$ for all $h \in \text{Hist}(\sigma) \cap \text{Hist}(\sigma')$;

- (2) $\text{Hist}_0(\sigma) = \text{Hist}_0(\sigma')$ and $\sigma(h) \simeq_h^* \sigma'(h)$ for all $h \in \text{Hist}_0(\sigma)$;
(3) $\sigma \approx^* \sigma'$

Proof. **Proof of (1) \Rightarrow (2) :** By induction and application of Lemma 14.

Proof of (2) \Rightarrow (3) We assume that $\forall h \in \text{Hist}_0(\sigma) = \text{Hist}_0(\sigma')$, $\sigma(h) \simeq_h^A \sigma'(h)$ and show that $\sigma \approx^A \sigma'$.

It suffices to take an arbitrary $\tau \in \Gamma_{-p}$ and show that $\mathbb{P}_{(\sigma,\tau)}(\Phi(p)) = 1$ iff $\mathbb{P}_{(\sigma',\tau)}(\Phi(p)) = 1$. We first show that:

$$\forall h \in \text{Hist}_0(\sigma) = \text{Hist}_0(\sigma'), \quad \mathbb{P}_{(\sigma,\tau)}(\text{Cyl}(h)) = \mathbb{P}_{(\sigma',\tau)}(\text{Cyl}(h)). \quad (3)$$

This can be shown by induction using Lemma 14 where the base case is $\mathbb{P}_{(\sigma,\tau)}(\text{Cyl}(s_{\text{init}})) = 1 = \mathbb{P}_{(\sigma',\tau)}(\text{Cyl}(s_{\text{init}}))$ and the induction step is that $\mathbb{P}_{(\sigma,\tau)}(\text{Cyl}(h)) = \mathbb{P}_{(\sigma',\tau)}(\text{Cyl}(h))$ implies:

$$\begin{aligned} \mathbb{P}_{(\sigma,\tau)}(\text{Cyl}(h's')) &= \mathbb{P}_{(\sigma,\tau)}(\text{Cyl}(h'))\delta_r(s, (\sigma(h'), \tau(h')))(s') \\ &= \mathbb{P}_{(\sigma',\tau)}(\text{Cyl}(h'))\delta_r(s, (\sigma'(h'), \tau(h')))(s') \\ &= \mathbb{P}_{(\sigma,\tau)}(\text{Cyl}(h's')) \end{aligned}$$

We decompose the set of runs as $\square\text{Val}_{p,0}^* \uplus \diamond\text{Val}_{p,1}^* \uplus \diamond\text{Val}_{p,-1}^*$. When intersecting with $\Phi(p)$, the last set in the union becomes empty and we have:

$$\Phi(p) = (\square\text{Val}_{p,0}^* \cap \Phi(p)) \uplus (\diamond\text{Val}_{p,1}^* \cap \Phi(p)).$$

Hence for $\sigma'' \in \{\sigma, \sigma'\}$, $\mathbb{P}_{(\sigma'',\tau)}(\Phi(p)) = \mathbb{P}_{(\sigma'',\tau)}(\square\text{Val}_{p,0}^* \cap \Phi(p)) + \mathbb{P}_{(\sigma'',\tau)}(\diamond\text{Val}_{p,1}^*)$ where we used that $\diamond\text{Val}_{p,1}^*$ implies $\Phi(p)$ almost surely because σ'' is winning from histories of value 1.

Hence it suffices to show that $\mathbb{P}_{(\sigma,\tau)}(\square\text{Val}_{p,0}^* \cap \Phi(p)) = \mathbb{P}_{(\sigma',\tau)}(\square\text{Val}_{p,0}^* \cap \Phi(p))$ and $\mathbb{P}_{(\sigma,\tau)}(\diamond\text{Val}_{p,1}^*) = \mathbb{P}_{(\sigma',\tau)}(\diamond\text{Val}_{p,1}^*)$. The former equality is due to (3). To prove the latter equality, we write for $\sigma'' \in \{\sigma, \sigma'\}$,

$$\mathbb{P}_{(\sigma'',\tau)}(\diamond\text{Val}_{p,1}^*) = \sum_{h \in \text{Hist}_0(\sigma'')} \mathbb{P}_{(\sigma'',\tau)}(\text{Cyl}(h))\delta_r(s, (\sigma''(h), \tau(h)))(\text{Val}_{p,1}^*). \quad (4)$$

This equality traduces the decomposition of the even $\text{Val}_{p,1}^*$ into the disjoint events parametrised by the history h of value 0 just before entering $\text{Val}_{p,1}^*$. Using (4), Lemma 14 and (3) we deduce that $\mathbb{P}_{(\sigma,\tau)}(\diamond\text{Val}_{p,1}^*) = \mathbb{P}_{(\sigma',\tau)}(\diamond\text{Val}_{p,1}^*)$, ending the proof.

Proof of (3) \Rightarrow (1): Straightforward. \blacktriangleleft

We are now ready to prove Theorem 8:

Proof of Theorem 8. We start with the proof of point (i).

" σ is \star -admissible" implies " σ is \star -SCO": We show the contrapositive. Let σ be a strategy that is not \star -SCO, and let us show that σ is not \star -admissible. Since σ is not \star -SCO, there is, by definition of \star -SCO, a history h s.t. $\chi_\sigma^*(h) \neq \chi_p^*(h)$. Recall that $\chi_p^*(h) = \max_{\sigma'} \chi_{\sigma'}^*(h)$. Hence, $\chi_\sigma^*(h) \neq \chi_p^*(h)$ implies that $\chi_p^*(h) \mathcal{G}\sigma < \chi_p^*(h)$ and that there exists a strategy σ' s.t. $\chi_{\sigma'}^*(h) = \chi_p^*(h)$. Consider the strategy $\bar{\sigma} = \sigma \langle h \leftarrow \sigma' \rangle$ that plays like σ and switches to σ' after history h . We claim that $\sigma \prec^* \bar{\sigma}$. Observe that, by construction: $\chi_\sigma^*(h) < \chi_{\bar{\sigma}}^*(h)$. Note also that $\chi_\sigma^h(\sigma) - 1$ is not possible, because this means that all continuations of h are losing for p , so it is not possible to have $\bar{\sigma}$ with $\chi_\sigma^*(h) < \chi_{\bar{\sigma}}^*(h)$. So we have necessarily $\chi_\sigma^h(\sigma)0$ and $\chi_\sigma^h(\bar{\sigma})1$. This means that $\bar{\sigma}$ is winning from h (against all possible strategies of $-p$), while, by definition of $\chi_\sigma^*(h) = 0$, σ is, from h , losing for p against at least one strategy of $-p$. Hence, $\sigma \prec^* \bar{\sigma}$ and σ is not \star -admissible.

" σ is \star -admissible" implies " σ is \star -LA": We show the contrapositive: we assume that σ is not \star -LA and show that σ is not \star -admissible. There exists a history $h \in \text{Hist}(\sigma)$ such that $\sigma(h)$ is \star -dominated at h by an \star -LA move b that can be chosen Dirac by virtue of Lemma 5 (i). There exists $a \in \text{Supp}(\sigma(h))$ such that $a <_h^* b$ (otherwise all the moves of the support of $\text{Supp}(\sigma(h))$ would be equivalent to the \star -LA move b and so would be α by Lemma 5 (ii)). We already saw in the proof of Lemma 5 (i), that $\sigma \preceq^* \sigma_a$ where σ_a is the strategy that plays like σ everywhere except in h where it plays a instead of α . Then it suffices to show that σ_a is \star -dominated. Using Lemma 7 and notation therein with $a \leq_h^* b$, we know that for every $c \in \Sigma_{-p}(s)$, $\chi_p^*(hs_{(a,c)}) \leq \chi_p^*(hs_{(b,c)})$ and if $\chi_p^*(hs_{(a,c)}) = \chi_p^*(hs_{(b,c)}) = 0$ then $s_{(a,c)} = s_{(b,c)}$. Here we can reuse the end of the proof of Lemma 7 to find a strategy σ' such that $\sigma_a \preceq^* \sigma'$ and that satisfies the further properties that it is \star -SCO at $hs_{(b,c)}$ for every $c \in \Sigma_{-p}(s)$, and it plays as h in every proper prefix of h . It remains to show that there exists a profile $\tau \in \Gamma_{-p}$ such that $(\sigma_a, \tau) \not\models^* \Phi(p)$ and $(\sigma', \tau) \models^* \Phi(p)$. Using Lemma 7, this time with $\sigma_a(h) \not\leq_h^* a$ we get the existence of a Dirac move $c \in \Sigma_{-p}(s)$ such that $\chi_p^*(hs_{(a,c)}) < \chi_p^*(hs_{(b,c)})$. We define $\tau \in \Gamma_{-p}$ such that $h \in \text{Hist}(\sigma_a, \tau)$, $\tau(h) = c$ and a further property depending on cases described as follows. If $\chi_{hs_a}^p = -1$ then $\chi_{hs_b}^p \geq 0$ and we further require that $(\sigma', \tau) \models_{hs_b}^* \Phi(p)$. If $\chi_{hs_a}^p = 0$ then $\chi_{hs_b}^p = 1$ and we further require that $(\sigma_a, \tau) \not\models_{hs_a}^* \Phi(p)$. In both cases it holds that $(\sigma', \tau) \models_{hs_b}^* \Phi(p)$ and $(\sigma_a, \tau) \not\models_{hs_a}^* \Phi(p)$. So $(\sigma_a, \tau) \not\models^* \Phi(p)$ and $(\sigma', \tau) \models^* \Phi(p)$. We conclude that σ is not \star -LA because $\sigma \preceq^* \sigma_a \prec^* \sigma'$.

" σ is \star -SCO and \star -LA" implies " σ is \star -admissible": Let σ be a strategy that is both \star -LA and \star -SCO. We show that for every σ' such that $\sigma \preceq^* \sigma'$ implies $\sigma \approx^* \sigma'$; which is a way of proving that σ is \star -Admissible. Let $\sigma' \in \Gamma_p$ such that $\sigma \preceq^* \sigma'$. We can assume without loss of generality that σ' is winning from histories of value 1 (otherwise it suffices to change σ' so that when a history of value 1 it switches to a winning strategy from that history). $\sigma \preceq^* \sigma'$ we deduce that $\sigma(h) \leq_h^* \sigma'(h)$ for every $h \in \text{Hist}(\sigma) \cap \text{Hist}(\sigma')$, then $\sigma(h) \simeq_h \sigma'(h)$ because $\sigma(h)$ is \star -LA at h . We can apply Lemma 15, and deduce that $\sigma \preceq^* \sigma'$ as required.

In simple safety games: " σ is \star -LA" implies " σ is \star -admissible": To establish this case, we prove that, in simple safety games, " σ is \star -LA" implies " σ is \star -SCO". This implies that all \star -LA strategies are also \star -SCO, hence admissible. First, observe that, in the case of simple safety games, the two semantics **A** and **S** coincide. Indeed, a winning strategy for the **S** semantics is also winning for the almost sure semantics. On the other hand, a winning strategy for the almost sure semantics ensures that no prefix generated by the strategy reaches the bad states (otherwise, this prefix would have non-negative measure, and the probability of winning would be < 1). Hence, a winning strategy for the almost sure semantics is also winning for the sure semantics. So, we can restrict ourselves to the sure semantics in the arguments that follow.

From σ , we build a new Dirac \star -LA strategy σ' , s.t., for all histories h , $\sigma'(h)$ is an arbitrarily chosen action from $\text{Supp}(\sigma(h))$. By Lemma 5, point (iii), σ' is also \star -LA. We claim that σ' is *value-preserving* i.e., for all histories h : $\chi_p^S(h) \leq \chi_p^S(\sigma'(h))$, for all $s \in \text{Succ}(\text{last}(h), \sigma'(h))$ (assuming σ , and hence σ' , are player p strategies). In other words, by playing σ' , player p never decreases the value of the histories he visits. This value preserving property holds by Lemma 7.

Moreover, since the objective is a simple safety one, it is easy to see that when σ' is **S**-winning, then σ is **S**-winning too. So σ' is dominated by σ and it is sufficient to prove that σ' is admissible to show that σ is admissible too.

So, when states of value 1 are entered, σ' ensures that they are never left, and σ' is thus winning from states of value 1. On the other hand, when a history of value 0 is entered, σ'

will also ensure that only histories of value 0 or 1 are visited. In both cases, the σ' is winning because we are considering a *safety objective*. In particular, if we stay in histories of value 0 forever, the strategy is winning because the bad states are never visited (otherwise, the value of the history would drop to -1).

We conclude that σ' is \star -SCO. Since it is also \star -LA, σ' is admissible. Hence, σ is admissible too.

Now, we move to the proof of point (ii).

Let σ be a strategy for player p , we build an admissible strategy that weakly dominates σ . This proof is similar to the proof of a similar result for SCO strategies in turn based games ([6, Lemma 7]). The strategy σ' is built dynamically along runs. It plays a winning strategy as soon as the current history begins to be of value 1. When the history begins to be of value -1 , there is nothing more to do, the strategy can play arbitrarily. Consider now, runs in which the value is always 0. The strategy σ' keeps in memory a wished strategy σ_h for himself and a wished strategy τ_h of player $-p$ such that $(\sigma_h, \tau_h) \models_h^* \Phi(p)$. The strategies σ_h and τ_h are inductively defined as follows. At the beginning, $h = s_{\text{init}}$, σ_h and τ_h are chosen such that $(\sigma_h, \tau_h) \models_h^* \Phi(p)$ and such that σ_h is h -admissible. Here and below if σ satisfies these properties then the choice $\sigma_h = \sigma$ is made by default, otherwise another strategy σ_h is chosen (and it dominates σ). After history h , σ' plays $\sigma_h(h)$. If after some history h the continuation hs' chosen is not in $\text{Hist}(\sigma_h, \tau_h)$ (this is the case only if player $-p$ plays something else than $\tau_h(h)$), then $\sigma_{hs'}$ and $\tau_{hs'}$ are chosen such that $\sigma_{hs'}$ is hs' -admissible and $(\sigma_{hs'}, \tau_{hs'}) \models_{hs'}^* \Phi(p)$. Otherwise, if the wished profile is followed then the wished strategy for player $-p$ is left unchanged ($\tau_{hs'} = \tau_h$); the wished strategy for player p is left unchanged ($\sigma_{hs'} = \sigma_h$) if σ_h is hs' -admissible and otherwise $\sigma_{hs'}$ is defined as a strategy that plays an admissible moves in hs' and that dominates σ_h . The constructed strategy is \star -LA. It is also \star -SCO because from every history h of value 0, $(\sigma', \tau_h) \models_h^* \Phi(p)$ and from every history of value 1, σ' is winning. σ' plays as σ by default or a strategy that dominates σ . At the end σ' is an admissible strategy that weakly dominates σ . ◀

B.5 Proof of Proposition 10

Proof. Take σ is \mathbf{A} -admissible. We define σ' as follows. From histories of value -1 , plays arbitrary Dirac strategies. From histories of value 1, plays winning strategies. If $\star = \mathbf{S}$ we can further requires that such winning strategies are Dirac. If there exists histories of value 0 then necessarily the initial state s_{init} is such a history and we build σ' inductively as follows such that $\text{Hist}_0(\sigma') = \text{Hist}_0(\sigma)$. Consider $h \in \text{Hist}_0(\sigma)$, then define $\sigma'(h)$ to be an arbitrary Dirac move $a \in \text{Supp}(\sigma(h))$. Then by Lemma 5 (iii) and the fact that $\sigma(h)$ is \star -LA we deduce that $\sigma'(h) \simeq_h^* \sigma(h)$; The histories of the form hs' that are in $\text{Hist}_0(\sigma')$ are exactly those that belongs to $\text{Hist}_0(\sigma)$ (see Lemma 14). By induction, $\text{Hist}_0(\sigma') = \text{Hist}_0(\sigma)$ and for every $h \in \text{Hist}_0(\sigma)$, $\sigma'(h) \simeq_h^* \sigma(h)$. We can use Lemma 15 and deduce that $\sigma \simeq_h^* \sigma'$ with σ' Dirac in histories of value ≤ 0 , and ≤ 1 when $\star = \mathbf{A}$. ◀

C Proofs of Section 4

C.1 Proof of Theorem 11

Before we prove the Theorem, we need several auxiliary lemmas.

► **Lemma 16.** *If a strategy $\sigma \in \Gamma_p(\mathcal{G})$ is \star -admissible then for all of its realisations $\bar{\sigma} \in \Gamma_p^{\text{det}}(\mathcal{G}_p^*), \mathcal{G}_p^*, \bar{\sigma} \models^S \Phi_0^*(p)$.*

Proof. Assume toward a contradiction that there is a \star -admissible strategy $\sigma \in \Gamma_p(\mathcal{G})$ for which there exists a realisation $\bar{\sigma} \in \Gamma_p^{\text{det}}(\mathcal{G}_p^*)$ that does not satisfy $\Phi_0^*(p)$. We can further assume wlog. that σ plays Dirac moves in states of value 0. There is a run $\rho \in \text{Outcome}(\mathcal{G}_p^*, \bar{\sigma})$ that does not satisfy $\Phi_0^*(p)$. More precisely, ρ satisfies

$$\Box \text{Val}_{i,0}^* \wedge \neg \Phi(i) \wedge \Diamond \Box \neg \text{AfterHelpMove}_p^*.$$

The run ρ does not satisfy $\Phi(p)$ but all values of its prefixes are equal to 0 and there is k_0 such that for every $k \geq k_0$, $\rho_k \notin \text{AfterHelpMove}_p^*$. We now show that for every $\rho' \in \text{Outcome}(\mathcal{G}, \sigma)$, such that $\mathfrak{o}(\rho_{\leq k_0}) \subseteq_{\text{pref}} \rho'$, it holds that $\rho' \notin \Phi(p)$. If $\rho' = \mathfrak{o}(\rho)$ then $\rho' \notin \Phi(p)$. If $\rho' \neq \mathfrak{o}(\rho)$ then there exists $k \geq k_0$ for which, $\rho'_{\leq k} = \mathfrak{o}(\rho)_{\leq k}$ and $\rho'_{\leq k+1} \neq \mathfrak{o}(\rho)_{\leq k+1}$. Let $h' = \rho'_{\leq k} = \mathfrak{o}(\rho)_{\leq k}$, $s = \text{last}(h')$, $a = \bar{\sigma}(\rho_{\leq k})$. h' as a prefix of $\mathfrak{o}(\rho)$ has value 0 so $\sigma(h')$ is Dirac and thus equal to a . Let $c, c' \in \Sigma_{-p}(s)$ such that $\delta(s, (a, c)) = \mathfrak{o}(\rho)_{k+1}$, $\delta(s, (a, c')) = \rho'_{k+1}$. From $\mathfrak{o}(\rho)_{k+1} \notin \text{AfterHelpMove}_p^*$ we know that $\chi_p^*(\rho'_{\leq k+1}) = \chi_p^*(h' \delta(s, (a, c'))) = -1$. We have found a prefix of ρ' that has value -1 , thus $\rho' \notin \Phi(p)$. We have proved that for all run $\rho' \in \text{Outcome}(\mathcal{G}, \sigma)$ such that $\mathfrak{o}(h) \subseteq_{\text{pref}} \rho'$, it holds that $\rho' \notin \Phi(p)$. This implies that $\chi_\sigma^*(\mathfrak{o}(h)) = -1 < 0 = \chi_p^*(\mathfrak{o}(h))$ proving that σ is not \star -SCO and hence not \star -admissible. This is a contradiction. \blacktriangleleft

► Lemma 17. *Given a \star -admissible strategy $\sigma \in \Gamma_p(\mathcal{G})$ and $\bar{\sigma} \in \Gamma_p^{\text{det}}(\mathcal{G}_p^*)$ one of its realisation. If a run $\rho \in \text{Outcome}(\mathcal{G}_p^*, \bar{\sigma})$ satisfies $\Box(\neg \text{Val}_{p,1}^*) \wedge \bigwedge_{q \neq p} \Phi_0^*(q) \wedge \Phi_1^*(q)$ in \mathcal{G}_p^* then there exists a profile $\tau \in \Gamma_{-p}(\mathcal{G})$ of \star -admissible strategies such that $\text{Outcome}(\mathcal{G}, (\sigma, \tau)) = \{\mathfrak{o}(\rho)\}$.*

Proof. For every $k \geq 1$ there exists moves $a_1^k, \dots, a_n^k \in \bar{\Sigma}$ such that $\delta(\rho_k, a_1^k, \dots, a_n^k) = \rho_{k+1}$ and $a_p^k = \bar{\sigma}(\rho_{\leq k})$. For every $q \neq p$, let $\tau_q \in \Gamma_q(\mathcal{G})$ be a strategy such that for every $k \geq 1$, $\tau_q(\mathfrak{o}(\rho_{\leq k})) = a_q^k$ and that plays an admissible strategy as soon as the current history is not a prefix of $\mathfrak{o}(\rho)$. We now show that $\text{Outcome}(\mathcal{G}, (\sigma, \tau)) = \{\mathfrak{o}(\rho)\}$. Denote by $a_{-p}^k = (a_q^k)_{q \in P \setminus \{p\}}$ the Dirac move played by $-p$ in $\mathfrak{o}(\rho_{\leq k})$. Since σ is \star -admissible; for every $k \geq 1$, $\sigma(\rho_{\leq k})$ is \star -LA and by Lemma 5 (iii) equivalent to the move of its support a_p^k , then by Lemma 14, we get $\delta_r(\mathfrak{o}(\rho_k), (\sigma(\mathfrak{o}(\rho_{\leq k})), a_{-p}^k))(\mathfrak{o}(\rho_{k+1})) = \delta_r(\mathfrak{o}(\rho_k), (a_p^k, a_{-p}^k))(\mathfrak{o}(\rho_{k+1})) = 1$. This implies that $\text{Outcome}(\mathcal{G}, (\sigma, \tau)) = \{\mathfrak{o}(\rho)\}$.

We now show that τ_q is admissible in \mathcal{G} for every $q \neq p$. It is clearly \star -LA as the moves a_q^k are \star -LAs and τ_q plays \star -LA moves at histories not prefix of $\mathfrak{o}(\rho)$. It is also \star -SCO at histories not prefix of $\mathfrak{o}(\rho)$. We now show that τ_q is \star -SCO at every $h \subseteq_{\text{pref}} \mathfrak{o}(\rho)$.

If $\chi_q^*(h) = 1$ then by $\Phi_1^*(q)$, ρ is winning for $\Phi(q)$, and hence so is $\mathfrak{o}(\rho)$. As a_q^k is a \star -LA move for every k , all histories of the form $\mathfrak{o}(\rho_{\leq k})$ for $k \geq |h|$ and their successors through a_q^k have value 1. Consider a profile σ such that $\sigma_q = \tau_q$, $h \in \text{Hist}(\sigma)$, we show that $\mathcal{G}, \sigma \models_h^* \Phi(q)$. We do this part of proof for $\star = \mathbf{A}$ (the proof for $\star = \mathbf{S}$ is a bit easier and omitted). We show that $\mathbb{P}_\sigma(\neg \Phi(q) \cap \text{Cyl}(h)) = 0$. For this, we decompose the set $\neg \Phi(q) \cap \text{Cyl}(h)$ into disjoint sets of runs:

$$\neg \Phi(q) \cap \text{Cyl}(h) = (\neg \Phi(q) \cap \{\mathfrak{o}(\rho)\}) \cup \bigcup_{k \geq |h|} \bigcup_{s'} (\neg \Phi(q) \cap \text{Cyl}(\mathfrak{o}(\rho_{\leq k})s')) \quad (5)$$

where the last union range over the state $s' \in \text{Supp}(\delta_r(\rho_k, \sigma(\mathfrak{o}(\rho_{\leq k}))) \setminus \{\mathfrak{o}(\rho_k)\})$. The histories $\mathfrak{o}(\rho_{\leq k})s'$ quantified above have value 1, and τ_q is \mathbf{A} -admissible at these histories by construction since the histories are not prefixes of $\mathfrak{o}(\rho)$. Hence, \mathbf{A} -winning from these histories, so $\mathbb{P}_\sigma(\neg \Phi(q) \cap \text{Cyl}(\mathfrak{o}(\rho_{\leq k})s')) = 0$. $\mathbb{P}_\sigma(\neg \Phi(q) \cap \{\mathfrak{o}(\rho)\}) = 0$ because $\mathfrak{o}(\rho) \in \Phi(q)$. The set $\neg \Phi(q) \cap \text{Cyl}(h)$ is a countable union of sets of probability 0, it is thus also of probability 0. The proof for the sure semantics is similar: it suffices to replace the fact that a set has probability zero by emptiness of this set. We have thus proved that $\chi_\sigma^*(h) = \chi_q^*(h) = 1$.

If $\chi_q^*(h) = 0$. We show that there exists a profile σ such that $\sigma_q = \tau_q$, $h \in \text{Hist}(\mathcal{G}, \sigma)$ and $\sigma \models^* \Phi(q)$. Let $k_0 = |h|$. If $\sigma(\rho) \in \Phi(q)$ then it suffices to take $\sigma = (\sigma, \tau)$. Otherwise we begin to show that there is $k \geq k_0$ such that $\rho_{k+1} \in \text{AfterHelpMove}_q^*$. Assume toward a contradiction that such a k does not exist, then $\Phi_0^*(q)$ tells us that there exists k_2 such that $\rho_{k_2} \in \text{Val}_{q,-1}^*$ or $\rho_{k_2} \in \text{Val}_{q,1}^*$. Necessarily $k_2 > k_0$ because as \star -LA moves only are played the states of value 1 cannot be escaped once entered. In the case $\rho_{k_2} \in \text{Val}_{q,1}^*$, ρ satisfies $\diamond \text{Val}_{q,1}^*$ and hence $\Phi(q)$ by virtue of $\Phi_1^*(q)$. This contradicts the fact that we are considering the case $\sigma(\rho) \notin \Phi(q)$. Consider now that $\rho_{k_2} \in \text{Val}_{q,-1}^*$. Let k_1 be the unique number in $\{k_0, \dots, k_2 - 1\}$ such that $\chi_q^*(\rho_{k_1}) = 0$ and $\chi_q^*(\rho_{k_1+1}) = -1$. Since $k_1 \geq k_0$, $\rho_{k_1+1} \notin \text{AfterHelpMove}_q^*$, from which we deduce that $\text{Succ}(\rho_{k_1}, a_q^k) \subseteq \text{Val}_{q,-1}^*$. This contradicts the fact that a_q^k is a \star -LA move in a history of value 0. We have showed the existence of a prefix length $k \geq |h|$ such that $\rho_{k+1} \in \text{AfterHelpMove}_{\rho_k, q}^*$ and we resume the proof of $\chi_q^*(h) = 0 = \chi_q^*(h)$. There exists a state $s'' \in \text{Succ}(a_q^{k+1}) \setminus \{\sigma(\rho_{k+1})\}$ such that $\chi_q^*(s'') \geq 0$. Moreover τ_q is \star -admissible from $\sigma(\rho_{\leq k})s''$, so there is a profile σ such that $\sigma(\rho_{\leq k})s'' \in \text{Hist}(\sigma)$, $\sigma_q = \tau_q$ and $\mathcal{G}, \sigma \models^* \Phi(q)$. Since $h \subseteq_{\text{pref}} \sigma(\rho_{\leq k})s''$, this implies that $\chi_\sigma^*(h) \geq 0$ as required.

It remains the case of histories h prefix of $\sigma(\rho)$ such that $\chi_q^*(h) = -1$. This case is straightforward because the following sequence of inequalities holds for every history $-1 \leq \chi_\sigma^*(h) \leq \chi_q^*(h)$. At the end, for every $q \neq p$, τ_q is \star -admissible as both \star -LA and \star -SCO, and $\sigma(\rho) = \text{Outcome}(\mathcal{G}, (\sigma, \tau))$. \blacktriangleleft

Before proving the main theorem of Section 4, we need to establish the following lemma:

► Lemma 18 (Admissible and dominant synthesis). *In concurrent games with prefix independent ω -regular winning conditions with semantics $\star \in \{\mathcal{S}, \mathcal{A}\}$, given a finite memory strategy σ for player p that is \star -LA, one can construct effectively a finite memory strategy σ' such that $\sigma \preceq^* \sigma'$.*

Proof. We assume that the strategy σ is given by a stochastic Moore machine \mathcal{M}_σ with set of memory states $M = \{m_1, m_2, \dots, m_k\}$, with a designated initial memory state m_{init} and such that after an history h , the machine is in memory state $\mathcal{M}_\sigma(h) \in M$. Each such memory state is labelled with a move noted $\mathcal{L}_\sigma(m)$. If we take the product of \mathcal{M}_σ with the game \mathcal{G} , we obtain a structure with states that are pairs (s, m) where s is a state of the original game and m is a memory state of \mathcal{M}_σ . This product can be seen as a new Moore machine that behaves exactly as \mathcal{M}_σ . From now on, we make the hypothesis that \mathcal{M}_σ has this state space.

We note $\chi_p^*(s)$ the value of state s in \mathcal{G} , this value is well defined as $\Phi(p)$ is prefix independent, and we note $\chi_\sigma^*(s, m)$ the value of the strategy σ for any history h that ends up in state s of the \mathcal{G} and memory state $m = \mathcal{M}_\sigma(h)$ of the machine.

As the strategy σ is \star -LA, if it is not admissible then it must be the case that there are state s , memory state m , and histories h such that after history h , the game \mathcal{G} is in state s , \mathcal{M}_σ is in state m , and $\chi_p^*(s) \neq \chi_\sigma^*(s, m)$, i.e. h is a witness that shows that σ is not SCO. Then, we will modify systematically \mathcal{M}_σ in a way that it behaves as \mathcal{M}_σ in all other histories and plays SCO in the histories that witness the fact that σ is not SCO.

We only need to consider the two following cases:

1. s and (s, m) are such that: $\chi_p^*(s) = 1$ and $\chi_\sigma^*(s, m) \leq 0$, then we replace (s, m) by a sub-machine that implements a \star -winning strategy from (s, m) . We know that finite memory machines that are \star -winning always exist and can be computed effectively for all ω -regular objectives.

2. s and (s, m) are such that: $\chi_p^*(s) = 0$ and $\chi_\sigma^*(s, m) = -1$, then we replace (s, m) by a sub-machine that plays any admissible strategy from s . We show this by establishing in the next claim that there is always such a finite state strategy and it can be computed effectively.

Claim. Given a game \mathcal{G} , a state s , and a player p , we can construct a finite state stochastic Moore machine that encodes an admissible strategy.

We establish this claim constructively. We consider the following case study to describe the machine:

- if the value of $\chi_p^*(s) = 1$, then the machine plays a finite state \star -winning strategy, such a finite memory strategy always exists and can be computed effectively.
- if the value of $\chi_p^*(s) = -1$, then the machine plays arbitrarily.
- if the value of $\chi_p^*(s) = 0$, then the machine selects a finite lasso-shape path $\rho = \rho_1 \cdot \rho_2^\omega$ such that ρ is compatible with \star -LA moves of player p and such that $\rho \in \Phi(p)$. Such a finite lasso-path always exists as $\chi_p^*(s) = 0$. Then the machine plays according to this lasso-path either forever or up to a deviation by another player. If the lasso path is played forever then the outcome is $\rho \in \Phi(p)$, and if there is a deviation, the new state of the game is s' and the three rules here are applied from s' (according to the value of state s').

Clearly, if entering a state s with value 0, we always choose the same finite lasso-path then the machine has finite state. As in states with value 0, it always plays \star -LA moves and ρ from s is such that $\rho \in \Phi(p)$, then we conclude that the strategy is \star -SCO and \star -LA and thus \star -admissible. So, we are done. ◀

We are now ready to Prove Theorem 11:

Proof of Theorem 11. Completeness: Assume there exists an admissible strategy $\sigma \in \Gamma_p(\mathcal{G})$ that wins against every admissible strategy. Let $\bar{\sigma} \in \Gamma_p^{\text{det}}(\mathcal{G})$ be a realisation of σ . Note that every runs $\rho \in \text{Outcome}(\mathcal{G}_p^*, \bar{\sigma})$ that satisfies $\diamond \text{Val}_{p,1}^*$ also satisfies $\Phi_{\mathcal{G}_p^*}$. The other runs $\rho \in \text{Outcome}(\mathcal{G}_p^*, \bar{\sigma})$ satisfy $\neg \diamond \text{Val}_{p,1}^* \equiv \square(\neg \text{Val}_{p,1}^*)$; for these runs we show that $\bigwedge_{q \neq p} \Phi_0^*(q) \wedge \Phi_1^*(q) \rightarrow \Phi(p)$. So let $\rho \in \text{Outcome}(\mathcal{G}_p^*, \bar{\sigma})$ be such that $\square(\neg \text{Val}_{p,1}^*) \wedge \bigwedge_{q \neq p} \Phi_0^*(q) \wedge \Phi_1^*(q)$. By Lemma 17, there exists τ that only contains admissible profiles such that $\text{Outcome}(\sigma, \tau) = \{\sigma(\rho)\}$. By assumption (σ, τ) is winning for $\Phi(p)$, hence so is ρ .

Correctness: We show the following stronger statement. If a strategy σ is winning for $\Phi_{\mathcal{G}_p^*}$ in \mathcal{G}_p^* then one can construct a strategy $\hat{\sigma}$ admissible that wins against every profile of admissible strategies as follows: (i) take an extension σ' of σ ; (ii) modify σ' into σ'' that plays a winning strategy as soon as a history of value 1 is entered; (iii) use Lemma C.1 to design an admissible strategy $\hat{\sigma}$ that weakly dominates σ'' .

We do the proof for the almost-sure semantics only, that is, for $\star = \mathbf{A}$. The slightly simpler but similar proof for $\star = \mathbf{S}$ is omitted. We first show that if a strategy σ is winning for $\Phi_{\mathcal{G}_p^*}$ in $\mathcal{G}_p^{\mathbf{A}}$ then for every extension σ' of σ and every profile τ of admissible strategies, it holds that $\mathcal{G}, (\sigma', \tau) \models^{\mathbf{A}} \Phi(p) \vee \diamond \text{Val}_{p,1}^{\mathbf{A}}$.

Let τ be a profile of admissible strategies. Note that $\mathbb{P}_{(\sigma', \tau)}(\diamond \text{Val}_{p,1}^{\mathbf{A}} \vee \Phi(p)) = 1$ is equivalent to $\mathbb{P}_{(\sigma', \tau)}(\square(\neg \text{Val}_{p,1}^{\mathbf{A}}) \wedge \neg \Phi(p)) = 0$. For $q \neq p$, τ_q is admissible, so, winning from histories of value 1, hence $\mathbb{P}_{(\sigma', \tau)}(\Phi_1^{\mathbf{A}}(q)) = 1$ and we have $\mathbb{P}_{(\sigma', \tau)}(\square(\neg \text{Val}_{p,1}^{\mathbf{A}}) \wedge \neg \Phi(p)) = \mathbb{P}_{(\sigma', \tau)}(\bigwedge_{q \neq p} \Phi_1^{\mathbf{A}}(q) \wedge \square(\neg \text{Val}_{p,1}^{\mathbf{A}}) \wedge \neg \Phi(p))$. To show that this probability is null it suffices to show that every run $\rho \in \text{Outcome}(\mathcal{G}, (\sigma', \tau))$ that satisfies $\bigwedge_{q \neq p} \Phi_1^{\mathbf{A}}(q) \wedge \square(\neg \text{Val}_{p,1}^{\mathbf{A}})$ also satisfies $\Phi(p)$. Take such a run ρ , then there is a sequence of move profile \mathbf{a}^k such that $\mathbf{a}_q^k \in \text{Supp}(\tau_q(\rho_k))$ and $\rho_{k+1} = \delta(\rho_k, \mathbf{a}^k)$. Define the run $\bar{\rho}$ by $\bar{\rho}_k = (\rho_k, \mathbf{a}^k)$ for $k \geq 1$ and $\bar{\rho}_0 = s'_{\text{init}}$ and note that $\rho = \sigma(\bar{\rho})$. We also have that $\bar{\rho}$ is a run of $\text{Outcome}(\mathcal{G}_p^{\mathbf{A}}, \sigma)$ and for

$q \neq p$ of $\text{Outcome}(\mathcal{G}_p^A, \bar{\tau}_q)$ where $\bar{\tau}_q$ is any realisation of τ_q for which $a_q^k = \tau_q(\rho_k)$ for every $k \geq 1$. The run $\bar{\rho}$ satisfies $\bigwedge_{q \neq p} \Phi_1^A(q)$ by assumption on ρ and $\bigwedge_{q \neq p} \Phi_0^A(q)$ because all the τ_q are admissible (Lemma 16). Putting all the things together $\bar{\rho}$ satisfies $\Phi(p)$ and so does ρ . We conclude that σ' satisfies that for every profile τ of admissible strategies, it holds that $(\sigma', \tau) \models^A \Phi(p) \vee \Diamond \text{Val}_{p,1}^A$. By construction σ'' wins against every profile τ for which $(\sigma', \tau) \models^A \Diamond \text{Val}_{p,1}^A$. So σ'' is A-winning against every profile τ of A-admissible profile. Since $\sigma'' \preceq^A \hat{\sigma}$, we deduce that $\hat{\sigma}$ is A-winning against every A-admissible profile. ◀