On the Complexity of Branching Games with Regular Conditions*

Marcin Przybyłko[1] and Michał Skrzypczak[2]

1 **University of Warsaw, University of New Caledonia**
M.Przybylko@mimuw.edu.pl

2 **University of Warsaw**
M.Skrzypczak@mimuw.edu.pl

Abstract

Infinite duration games with regular conditions are one of the crucial tools in the areas of verification and synthesis. In this paper we consider a branching variant of such games—the game contains branching vertices that split the play into two independent sub-games. Thus, a play has the form of an infinite tree. The winner of the play is determined by a winning condition specified as a set of infinite trees. Games of this kind were used by Mio to provide a game semantics for the probabilistic μ-calculus. He used winning conditions defined in terms of parity games on trees. In this work we consider a more general class of winning conditions, namely those definable by finite automata on infinite trees. Our games can be seen as a branching-time variant of the stochastic games on graphs.

We address the question of determinacy of a branching game and the problem of computing the optimal game value for each of the players. We consider both the stochastic and non-stochastic variants of the games. The questions under consideration are parametrised by the family of strategies we allow: either mixed, behavioural, or pure.

We prove that in general, branching games are not determined under mixed strategies. This holds even for topologically simple winning conditions (differences of two open sets) and non-stochastic arenas. Nevertheless, we show that the games become determined under mixed strategies if we restrict the winning conditions to open sets of trees. We prove that the problem of comparing the game value to a rational threshold is undecidable for branching games with regular conditions in all non-trivial stochastic cases. In the non-stochastic cases we provide exact bounds on the complexity of the problem. The only case left open is the 0-player stochastic case, i.e. the problem of computing the measure of a given regular language of infinite trees.

1998 ACM Subject Classification F.4.3 Formal Languages

Keywords and phrases stochastic games, meta-parity games, infinite trees, regular languages, parity automata

Digital Object Identifier 10.4230/LIPIcs.MFCS.2016.79

1 Introduction

Since the seminal works of Büchi and Landweber [4], and of McNaughton [19], the infinite duration games are widely used to model interaction between a system and an environment. One of the fundamental questions about such games is the question of determinacy—does always one of the players has a winning strategy? In a more general case of valued zero-sum

* The authors were supported by the Polish National Science Centre grant no. 2014-13/B/ST6/03595.

41st International Symposium on Mathematical Foundations of Computer Science (MFCS 2016).
Editors: Piotr Faliszewski, Anca Muscholl, and Rolf Niedermeier; Article No. 79; pp. 79:1–79:25
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

games, determinacy amounts to the equality

$$\sup_{\sigma} \inf_{\pi} \; val(\sigma, \pi) = \inf_{\pi} \sup_{\sigma} \; val(\sigma, \pi), \tag{1}$$

where σ and π range over strategies of the respective players. It is often crucial to provide a specific information about the strategies that are enough to win a given game. Büchi and Landweber proved that if the winning condition of a game is a regular language of infinite words then the game is determined under finite memory strategies. Further results have established more precise bounds for the amount of memory needed [10, 11]. Also, the stochastic variant of the question was considered [8].

In this work we study a branching variant of stochastic games on graphs—a variant called *branching games* (also known as *tree games*, c.f. [22, Chapter 4]), played on *branching boards*. A play of a branching game consists of a number of threads, each thread develops independently. When a thread reaches a vertex marked as *branching*, it is split into two separate threads. Thus, a play of a branching game has the shape of an infinite tree. The winner of the play is determined by a winning condition specified as a regular set of infinite trees. Since the choices made by players in separate threads are independent, branching games are not games of perfect information, nor of perfect recall in the meaning of [3]. Games of this kind were used by Mio [23] to provide a game semantics for the probabilistic μ-calculus, with the *meta-parity* winning condition defined in terms of parity games on trees. The author called them *meta-parity games* and established their pure determinacy subject to some set-theoretic assumptions, that were later eliminated [15]. This result is interesting, because, as the author notices in conclusion, determinacy results for imperfect information games are not frequent in game theory.

In this article we address the question of determinacy of branching games and the problem of computing the game value for a more general class of winning conditions than the meta-parity conditions studied by Mio, namely those definable by finite automata on infinite trees. We believe that this extension is motivated by the role tree automata play in verification theory. Recall that these automata, introduced by Rabin in his proof of decidability of the Monadic Second-Order theory of the full k-ary tree [27], constitute a general formalism subsuming most of temporal logics of programs. The first step was made by the first author who extended the results of Mio [23] to winning conditions given by game automata [26].

We consider both the stochastic and non-stochastic variants of the games. The questions under consideration are parametrised by the family of strategies we allow, either mixed, behavioural, or pure. The goal of this work is to provide answers to two questions:

- When a branching game is determined in the sense of (1)?
- When the optimal value for a given player can be effectively computed?

Both questions can be asked for stochastic and non-stochastic variants of games, which usually yields different answers, see e.g. [8]; and for different sets of allowed strategies. The distinction between the sets of pure, behavioural, and mixed strategies can significantly alter the techniques and expected outcomes, see e.g. [6].

The answers we provide create an almost complete picture from the point of view of topological complexity of sets:

- non-stochastic branching games are not determined under pure nor behavioural strategies even for winning conditions that are topologically both closed and open,
- non-stochastic branching games are not determined under mixed strategies for winning conditions that are a difference of two open sets,
- non-stochastic branching games are determined under mixed strategies for winning conditions that are open (equivalently, closed) sets,

- the problem of comparing the value of a branching game to a rational number is undecidable in all the non-trivial stochastic cases,
- in the non-stochastic case, when we ask about the existence of a pure winning strategy, the problem is decidable and we provide precise bounds on its complexity.

The only remaining question is whether the value of a 0-player stochastic branching game can be effectively computed. This is equivalent to asking about computability of the measure of a given regular language of infinite trees.

Although the results of this paper show intractability of branching games, it is still possible that they are determined for a reasonable class of winning conditions. These ideas are discussed in Conclusions.

1.1 Related work

It is known that games with arbitrary pay-off functions are not determined. The celebrated result of Martin [17] states that *Gale-Steward games* with Borel payoffs are determined under pure strategies. His later result establishes an analogous result for the so-called *Blackwell games* and mixed strategies (cf. [18]). In this work we show that branching games are determined only for topologically simplest winning conditions—open or closed sets; even allowing a difference of two open sets leads to indeterminacy.

Since the branching games are not games of *perfect recall* the Kuhn's theorem (cf. e.g. [3]) does not hold; and the behavioural strategies are weaker than the mixed strategies. An example of such a situation was provided by Mio in [22, Chapter 4]. Therefore, there are three variants of the question of determinacy: pure, behavioural, and mixed determinacy.

The concept of branching games is a natural extension of the *meta-parity games* introduced by Mio [23] to provide a game semantics for the probabilistic μ-calculus. The first author proved in [26] that non-stochastic branching games with winning conditions given by the so-called *game automata* are determined under pure strategies. In this work we study determinacy of branching games from the perspective of the topological complexity of the winning condition.

Recently Asarin et al. considered the so-called *entropy games*, cf. [2], which can be easily embedded in our framework. In the authors' own words, an entropy game is played on a finite arena by two-and-a-half players: Despot, Tribune, and the non-deterministic People. The pay-off function is the entropy of the language formed by paths of the resulting tree. The authors of [2] prove that the entropy games are determined under pure strategies and can be solved in $NP \cap coNP$, extending the class of objectives for which branching games are determined.

The question of computing the coin-flipping measure of a given regular languages of infinite trees is one of the crucial open problems about probabilistic logics on infinite trees. Chen, Dräger, and Kiefer proved in [9] that the problem is decidable for regular languages recognisable by deterministic automata. The result was later strengthened by Michalewski and Mio [21] to the so-called game automata. The question of computing the value of a branching game is a natural extension of the above problem obtained by allowing interplay between the players. The first author implicitly provided bounds on the complexity of the problem in the non-stochastic case. In this paper we complete those bounds by proving *2-EXP*-hardness of the problem. Additionally, we prove that the problem becomes undecidable if any form of randomisation is allowed (either by considering randomised or behavioural strategies; or by adding stochastic positions). The only remaining open question is the original one—when there are only stochastic positions and no players.

Branching games fall into a general category of games of *imperfect information*, i.e. games where the full information about the state of the game is not assumed: the definitions assure that the players have no information about the execution of separate threads. The area of imperfect information games is rich and not fully understood, see e.g. [8, 5, 7], see also [6]. In this context, branching games with regular objectives can be seen as a natural extension of *imperfect information games with ω-regular objectives* to the branching-time case.

2 Definitions

In this section we will define the objects studied in the paper. The crucial definitions are those of a *branching game* and *game values*. By ω we denote the set of natural numbers and $\mathbb{R}$ stands for the set of reals.

Words and trees An *alphabet* Γ is a finite non-empty set. A *word* over Γ is any, possibly infinite, sequence $w = a_0 a_1 \cdots a_n \cdots$ where $a_i \in \Gamma$. By $w[i]$ we denote the i-th letter of w, i.e. a_i. ε stands for the empty sequence. Words are either *finite* (Γ^*) or *infinite* (Γ^ω). $|w|$ is the length of a finite word w. The prefix order on words is denoted $\sqsubseteq$.

A *tree* over an alphabet Γ is any partial function $t\colon \{\mathtt{L},\mathtt{R}\}^* \rightharpoonup \Gamma$ with a non-empty prefix-closed domain $\mathit{Dom}(t) \subseteq \{\mathtt{L},\mathtt{R}\}^*$. The elements $d \in \{\mathtt{L},\mathtt{R}\}$ are called *directions*, $\bar{d}$ is the direction opposite to d. Elements of the set $\{\mathtt{L},\mathtt{R}\}^*$ are called *nodes*. We say that a node u of a tree t is *fully branching* if it has two children in the tree, is *uniquely branching* if it has exactly one child in the tree. The set $\mathcal{T}_\Gamma$ is the set of all trees over an alphabet Γ. This set can naturally be enhanced with a topology in such a way that it becomes a homeomorphic copy of the Cantor set [29]. We say that a tree t_1 is a *prefix* of a tree t_2 if $t_1 \subseteq t_2$, i.e. $\mathit{Dom}(t_1) \subseteq \mathit{Dom}(t_2)$ and for every $u \in \mathit{Dom}(t_1)$ we have $t_1(u) = t_2(u)$.

Regular languages In this work we use the standard notions of non-deterministic and alternating parity automata over infinite trees. Together with Monadic Second-Order logic, these automata form equivalent formalisms for defining *regular* languages of infinite trees. For an introduction to this area see for instance [28].

Branching games This paper is about branching games. The two adversaries of our games are called Eve and Adam (or shortly E and A). Since we consider stochastic games, we additionally introduce *Nature* denoted $\mathcal{N}$. A *branching board* is a tuple $\mathtt{B} = \langle V, \Gamma, s_\mathtt{L}, s_\mathtt{R}, \rho, \eta, \lambda, v_\mathrm{I}\rangle$, where V is the set of *vertices*; Γ is the *alphabet*; $s_\mathtt{L}, s_\mathtt{R}\colon V \to V$ are the *successor functions*; $\lambda\colon V \to \Gamma$ is the *labelling* of the vertices; $\rho\colon V \to \{A, E, \mathcal{N}, \mathcal{B}\}$ is a partition of the vertices between Adam's, Eve's, *Nature*'s, and branching vertices; $\eta\colon \rho^{-1}\big(\{\mathcal{N}\}\big) \to \mathit{Dist}(\{\mathtt{L},\mathtt{R}\})$ maps *Nature*'s vertices to random distributions over the successors; $v_\mathrm{I} \in V$ is the *initial vertex*. We extend the assignment s to arbitrary sequences of directions in the natural way: $s_\varepsilon(v) = v$ and $s_{u\cdot d}(v) = s_d\big(s_u(v)\big)$.

For $P \in \{A, E, \mathcal{N}, \mathcal{B}\}$, by V_P we denote the set of vertices belonging to P, i.e. $\rho^{-1}\big(\{P\}\big)$. We say that $\mathtt{B}$ is *finitary* if the set of vertices V is finite and the values used to define η are rational. For $\mathcal{P} \subseteq \{A, E, \mathcal{N}, \mathcal{B}\}$ we say that $\mathtt{B}$ is $\mathcal{P}$-*branching* if $\mathit{Range}(\rho) \subseteq \mathcal{P}$. Every board $\mathtt{B}$ defines the tree $t^\lambda_\mathtt{B}\colon \{\mathtt{L},\mathtt{R}\}^* \to \Gamma$ as the unfolding of the adequate labelled sub-graph of the board, i.e. $t^\lambda_\mathtt{B}(u) = \lambda\big(s_u(v_\mathrm{I})\big)$. $\mathtt{B}$ is *non-stochastic* if it is $\{E, A, \mathcal{B}\}$-branching.

Intuitively, a play over a branching board $\mathtt{B}$ proceeds in threads, each thread has one token located in a vertex of the board. Initially, there is one thread with the token located in v_I. Consider a thread with a token located in a vertex v. If $\rho(v) = \mathcal{B}$ then the thread is duplicated into two separate threads with tokens located in $s_\mathtt{L}(v)$ and $s_\mathtt{R}(v)$. If $\rho(v) = \mathcal{N}$ then the token is moved either to $s_\mathtt{L}(v)$ or to $s_\mathtt{R}(v)$ depending on an independent random event with distribution $\eta(v)$. If $\rho(v) \in \{E, A\}$ then the respective player can make her/his choice depending on the history of the current thread. However, she/he cannot take into account positions of tokens from other threads in the current play. After all the threads moved infinitely many times, a tree-like play has been created. The winning condition of a branching game will indicate which plays are winning for which player. Figure 1 depicts a branching board and a play on this board.

We will now formalise the notions of a play and a pure strategy of a player. Consider a non-empty set $\mathcal{P} \subseteq \{A, E, \mathcal{N}, \mathcal{B}\}$. We say that a tree $t \subseteq t^\lambda_{\mathtt{B}}$ is *$\mathcal{P}$-branching* if it is fully branching in the nodes $u \in \{\mathtt{L}, \mathtt{R}\}^*$ such that $\rho\big(s_u(v_\mathrm{I})\big) \in \mathcal{P}$ and uniquely branching in the remaining nodes. A *play* on a board $\mathtt{B}$ is a tree $t \subseteq t^\lambda_{\mathtt{B}}$ that is $\{\mathcal{B}\}$-branching. The set of all plays on a board $\mathtt{B}$ is denoted $\mathrm{plays}(\mathtt{B})$. For $P \in \{E, A, \mathcal{N}\}$ we say that a tree $t \subseteq t^\lambda_{\mathtt{B}}$ is a *pure strategy* of P over $\mathtt{B}$ if t is $\big(\{E, A, \mathcal{N}, \mathcal{B}\} \setminus \{P\}\big)$-branching. The set of pure strategies of P over $\mathtt{B}$ is denoted $\Sigma^P_{\mathtt{B}}$. Notice that the sets $\mathrm{plays}(\mathtt{B})$ and $\Sigma^P_{\mathtt{B}}$ for $P \in \{E, A, \mathcal{N}\}$ are closed sets of Γ-labelled trees. If V is finite then all these sets are regular.

Given three pure strategies $\sigma \in \Sigma^E_{\mathtt{B}}$, $\pi \in \Sigma^A_{\mathtt{B}}$, and $\eta \in \Sigma^{\mathcal{N}}_{\mathtt{B}}$ the *play resulting from σ, π, and η* (denoted $\mathrm{eval}_{\mathtt{B}}(\sigma, \pi, \eta)$) is the tree $\sigma \cap \pi \cap \eta \in \mathrm{plays}(\mathtt{B})$. Thus, $\mathrm{eval}_{\mathtt{B}} \colon \Sigma^E_{\mathtt{B}} \times \Sigma^A_{\mathtt{B}} \times \Sigma^{\mathcal{N}}_{\mathtt{B}} \to \mathrm{plays}(\Gamma)$. Notice that the function $\mathrm{eval}_{\mathtt{B}}$ is continuous.

Measure theory For an introduction to measure theory we refer to [16, Chapter 17]. Measure properties of regular sets of trees are discussed in [15]. Let μ be a Borel measure on a topological space X. We say that μ is a *probability measure* if $\mu(X) = 1$. A function $f \colon X \to \mathbb{R}$ is μ-measurable if the pre-image of any measurable set in $\mathbb{R}$ is μ-measurable in X. $f \colon X \to \mathbb{R}$ is *universally measurable* if it is μ-measurable for every Borel measure μ on X. If $f \colon X \to \mathbb{R}$ is μ-measurable then by $\int_X f(x)\, \mu(\mathrm{d}x)$ we denote the integral of f with respect to the measure μ.

Branching games A *branching game* is a pair $G = \langle \mathtt{B}, \Phi \rangle$ where $\mathtt{B}$ is a branching board and Φ is a universally measurable bounded real function $\Phi \colon \mathrm{plays}(\mathtt{B}) \to \mathbb{R}^+$. The notions of a *$\mathcal{P}$-branching* game and a *finitary* game refer to the respective properties of the board.

Mixed strategies A mixed strategy of a player $P \in \{E, A\}$ is a Borel probability measure over the set $\Sigma^P_{\mathtt{B}}$. The set of all mixed strategies of P is denoted by $\Sigma^{MP}_{\mathtt{B}}$.

There is a natural way of defining a Borel probability measure $\eta^*_{\mathtt{B}}$ on the set $\Sigma^{\mathcal{N}}_{\mathtt{B}}$ of strategies of $\mathcal{N}$. This measure represents the intuition, that after a sequence of directions $u \in \{\mathtt{L}, \mathtt{R}\}$ corresponding to a vertex $v = s_u(v_\mathrm{I}) \in V$ such that $\rho(v) = \mathcal{N}$, *Nature* chooses to move to a direction $d \in \{\mathtt{L}, \mathtt{R}\}$ with the probability $\eta(v)(d)$ and is called *behavioural*.

Behavioural strategies We say that a mixed strategy of P is *behavioural* if it is a "coin flipping" measure, i.e. a measure induced by supplying some of the nodes of $t^\lambda_{\mathtt{B}}$ corresponding to vertices of P with a probability distribution over the successors. To produce a pure strategy from a behavioural one, the directions are chosen independently according to the fixed probability distributions.

More formally, a mixed strategy τ of P is *behavioural* if it is, as a measure over $\Sigma^P_{\mathtt{B}}$, the measure $\eta^*_{\mathtt{B}'}$ for some (possibly not finitary) board $\mathtt{B}'$. The set of all behavioural strategies of P is denoted by $\Sigma^{BP}_{\mathtt{B}}$. Clearly we can treat every pure strategy in $\Sigma^P_{\mathtt{B}}$ as a Dirac delta function in $\Sigma^{MP}_{\mathtt{B}}$ (in fact in $\Sigma^{BP}_{\mathtt{B}}$). Thus, we can assume that $\Sigma^P_{\mathtt{B}} \subseteq \Sigma^{BP}_{\mathtt{B}} \subseteq \Sigma^{MP}_{\mathtt{B}}$.

Strategies as functions There is a different way to define the three types of strategies that may give more intuition to the behaviour and expressive power of the strategies. A pure strategy $\sigma \in \Sigma^P$ can be seen as a function $\sigma \colon \{\mathtt{L}, \mathtt{R}\}^* \to \{\mathtt{L}, \mathtt{R}\}$; a behavioural strategy $\sigma_b \in \Sigma^{BP}$ as a function $\sigma_b \colon \{\mathtt{L}, \mathtt{R}\}^* \to \mu(\{\mathtt{L}, \mathtt{R}\})$; and a mixed strategy $\sigma_m \in \Sigma^{MP}$ as a measure

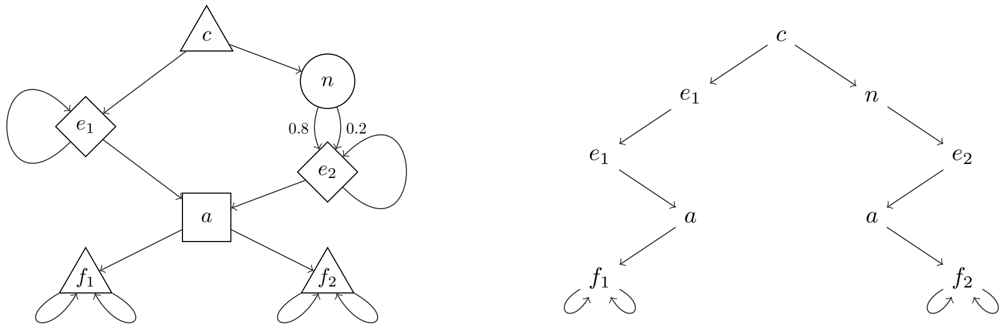

Figure 1 An example of a branching board and a play on this board. We denote Eve's, Adam's, *Nature*'s, and branching vertices by diamonds, squares, circles, and triangles respectively. *Nature*'s vertices are equipped with a probability distribution over the successors. The successors L and R agree with the directions on the picture, i.e. L moves to the left.

$\sigma_m \in \mu(\Sigma^P)$, where $\mu(X)$ dentotes some Borel probability mesure on the set X.

An example Figure 1 depicts a branching board B and a play t on this board. We identify the vertices with their labels. A pure strategy of Adam can make different choices in a depending on the history of the thread that lead to this vertex (there are infinitely many such histories). A pure strategy of Eve can make different choices in e_2 depending on the edge taken by *Nature* in n. A mixed strategy of Eve can synchronise: with probability $\frac{1}{2}$ move to L in both vertices e_1, e_2; and with probability $\frac{1}{2}$ move to R in both of them. A behavioural strategy cannot make such a synchronisation: the probability distribution over the successors depends only on the history of the current thread.

Values of strategies Assume that $\sigma_m \in \Sigma_{\mathtt{B}}^{ME}$ and $\pi_m \in \Sigma_{\mathtt{B}}^{MA}$ are two mixed strategies of the respective players. Our aim is to define the value $val_{\mathrm{G}}(\sigma_m, \pi_m)$. Intuitively, $val_{\mathrm{G}}(\sigma_m, \pi_m)$ should be the expected value of $\Phi\big(\mathrm{eval}_{\mathrm{G}}(\sigma, \pi, \eta)\big)$ where the pure strategies σ, π, and η are chosen according to the probability distributions σ_m, π_m, and $\eta^*_{\mathtt{B}}$ respectively. This is formalised as follows.

$$val_{\mathrm{G}}(\sigma_m, \pi_m) \stackrel{\text{def}}{=} \int_{\Sigma_{\mathtt{B}}^E, \Sigma_{\mathtt{B}}^A, \Sigma_{\mathtt{B}}^N} \Phi\big(\mathrm{eval}_{\mathrm{G}}(\sigma, \pi, \eta)\big)\ \sigma_m(\mathrm{d}\sigma)\, \pi_m(\mathrm{d}\pi)\, \eta^*_{\mathtt{B}}(\mathrm{d}\eta) \tag{2}$$

If σ and π are pure strategies and the board is non-stochastic then $val_{\mathrm{G}}(\sigma, \pi) = \Phi(\pi \cap \sigma)$.

Values of a game The aim of Eve in a branching game is to maximise the value $val_{\mathrm{G}}(\sigma, \pi)$. Let us define the *partial values* of the game. Consider $X \in \{\varepsilon, B, M\}$ (i.e. X stands for respectively *pure*, *behavioural*, and *mixed* strategies). The X *value of* G *for Eve* (resp. Adam) is defined as

$$val_{\mathrm{G}}^{XE} \stackrel{\text{def}}{=} \sup_{\sigma \in \Sigma_{\mathtt{B}}^{XE}} val_{\mathrm{G}}(\sigma) \quad \text{where} \quad val_{\mathrm{G}}(\sigma) \stackrel{\text{def}}{=} \inf_{\pi \in \Sigma_{\mathtt{B}}^{A}} val_{\mathrm{G}}(\sigma, \pi),$$

$$val_{\mathrm{G}}^{XA} \stackrel{\text{def}}{=} \inf_{\pi \in \Sigma_{\mathtt{B}}^{XA}} val_{\mathrm{G}}(\pi) \quad \text{where} \quad val_{\mathrm{G}}(\pi) \stackrel{\text{def}}{=} \sup_{\sigma \in \Sigma_{\mathtt{B}}^{E}} val_{\mathrm{G}}(\sigma, \pi).$$

Notice, that the second inf/sup is taken over the pure strategies of the opponent. This is explained by the following simple lemma.

▶ **Lemma 1.** *Let* G *be a branching game. If* σ_m *is Eve's mixed strategy then*

$$\inf_{\pi_m \in \Sigma_{\mathtt{B}}^{MA}} val_{\mathrm{G}}(\sigma_m, \pi_m) = \inf_{\pi_b \in \Sigma_{\mathtt{B}}^{BA}} val_{\mathrm{G}}(\sigma_m, \pi_b) = \inf_{\pi \in \Sigma_{\mathtt{B}}^{A}} val_{\mathrm{G}}(\sigma_m, \pi)$$

Dually, the same holds for mixed strategies of Adam if we replace inf *with* sup *and* A *with* E*.*

Determinacy As a simple consequence of Lemma 1 we obtain the following inequalities

$$val^{A}_{\mathrm{G}} \geq val^{BA}_{\mathrm{G}} \geq val^{MA}_{\mathrm{G}} \geq val^{ME}_{\mathrm{G}} \geq val^{BE}_{\mathrm{G}} \geq val^{E}_{\mathrm{G}}. \tag{3}$$

The first two (resp. the last two) inequalities hold by the fact that we take inf (resp. sup) over greater (reps. smaller) sets of strategies. The third inequality holds by Lemma 1 and the fact that $\inf_x \sup_y f(x,y) \geq \sup_y \inf_x f(x,y)$.

We will say that a branching game G is *determined*

- *under pure strategies* if $val^{A}_{\mathrm{G}} = val^{E}_{\mathrm{G}}$,
- *under behavioural strategies* if $val^{BA}_{\mathrm{G}} = val^{BE}_{\mathrm{G}}$,
- *under mixed strategies* if $val^{MA}_{\mathrm{G}} = val^{ME}_{\mathrm{G}}$.

Clearly, Equation (3) shows that pure determinacy implies behavioural determinacy and behavioural determinacy implies mixed determinacy. In general, the opposite implications do not hold. The questions of determinacy of branching games are discussed in Section 3.

Regular branching games The following theorem implies that we can take as Φ an indicator of a regular language of trees $L \subseteq \mathrm{plays}(\mathtt{B})$, i.e. $\Phi(t) = 1$ if $t \in L$ and $\Phi(t) = 0$ otherwise. In that case we say that a game G has L *as a winning condition* and we write $G = \langle \mathtt{B}, L\rangle$ instead of $G = \langle \mathtt{B}, \Phi\rangle$.

▶ **Theorem 2** (Michalewski et al. [15]). *Every regular language L of infinite trees is universally measurable, i.e. for every Borel measure μ on the set of trees, we know that L is μ-measurable.*

3 Determinacy

In this section we study determinacy of branching games in the three variants: pure, behavioural, and mixed; see (3). We will show that for general regular winning conditions all three variants fail. However, when we restrict to closed regular winning sets we can recover the mixed determinacy.

Notice that if a branching game is not branching, i.e. it is a $\{E, A, \mathcal{N}\}$-branching game then the determinacy is well-understood [18, 8]. Similarly, if there are no positions of one of the players then the game is purely determined by Lemma 1. Therefore, the simplest case specific for the branching games are the $\{E, A, \mathcal{B}\}$-branching games.

3.1 Behavioural indeterminacy

We start by proving the following theorem.

▶ **Theorem 3.** *There is a $\{E, A, \mathcal{B}\}$-branching game* G *with a regular winning condition that is both closed and open such that* G *is not determined under behavioural strategies.*

The board B of the game G is depicted in Figure 2. A play $t \in \mathrm{plays}(\mathtt{B})$ over B starts by splitting into four separate threads by the $\mathcal{B}$-vertices labelled with c. Then, each of the players can perform two separate choices, E in the two vertices labelled x_1 and x_2, and A in the two vertices labelled x_3 and x_4. Their choices lead to vertices labelled by either 0 or 1. The rest of the play stays forever in the branching vertex labelled by f. For $i = 1, 2, 3, 4$ let $x_i(t) \in \{0, 1\}$ be label chosen by the respective player in the vertex labelled by x_i, i.e. the label of the unique child of the unique node labelled by x_i in t. Consider a winning set $L \subseteq \mathrm{plays}(\mathtt{B})$ defined as follows

$$L \stackrel{\text{def}}{=} \big\{t \in \mathrm{plays}(\mathtt{B}) \mid x_1(t) = x_2(t) = x_3(t) = x_4(t) \ \vee\ x_3(t) \neq x_4(t)\big\} \tag{4}$$

In other words, Eve wins a play t if either Adam has chosen two different labels in x_3 and x_4 or all the chosen labels are equal. Since the vertices labelled x_i lie at a fixed depth of every play $t \in \mathrm{plays}(\mathtt{B})$, L is a closed and open regular language of infinite trees.

▶ **Example 4.** The game $\mathrm{G} = \langle \mathtt{B}, L\rangle$ has the following partial values:

$$val^{A}_{\mathrm{G}} = 1; \qquad val^{BA}_{\mathrm{G}} = \frac{3}{4}; \qquad val^{MA}_{\mathrm{G}} = \frac{1}{2} = val^{ME}_{\mathrm{G}}; \qquad val^{BE}_{\mathrm{G}} = \frac{1}{4}; \qquad val^{E}_{\mathrm{G}} = 0.$$

We first argue about the pure values—a pure strategy over the board from Figure 2 needs to declare in advance the two values $x_i(t)$ and $x_{i+1}(t)$ for $i = 1, 3$ depending on the player. If such a strategy is fixed, the opponent can choose his values in such a way to win.

We now consider the mixed value. Let σ_m randomly choose with equal probability between the following two pure strategies σ_i for $i = 0, 1$: the strategy σ_i satisfies $x_1(\sigma_i) = x_2(\sigma_i) = i$. π_m is defined analogously. It is easy to check that these strategies are optimal and witness that the mixed value of the game is $\frac{1}{2}$ for both players.

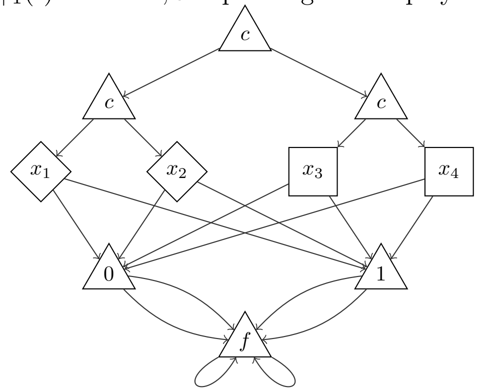

Figure 2 A branching board that is not determined under behavioural strategies.

Consider a behavioural strategy σ_b of Eve (the case of Adam is entirely dual). Such a strategy can be described by two independent random choices:

1. σ_b chooses x_1 to be 0 with probability p_1,
2. σ_b chooses x_2 to be 0 with probability p_2.

Thus, each behavioural strategy of Eve is characterised by a pair of numbers $p_1, p_2 \in [0, 1]$. A simple computation shows that no matter how Eve chooses her values p_1, p_2, Adam can find a counter-strategy guaranteeing the value of at most $\frac{1}{4}$.

Since $val^{BA}_{\mathrm{G}} = \frac{3}{4} \neq \frac{1}{4} = val^{BE}_{\mathrm{G}}$ the proof of Theorem 3 is concluded.

3.2 Mixed indeterminacy

We will now show that the mixed determinacy fails for relatively simple regular sets, as expressed by the following theorem.

▶ **Theorem 5.** *There is a $\{E, A, \mathcal{B}\}$-branching game with a regular winning set being a difference of two open sets that is not determined under mixed strategies.*

To prove this theorem we will encode the following game as a $\{E, A, \mathcal{B}\}$-branching game G. Assume that ∞ is an additional symbol such that for every $n \in \omega$ we have $n < \infty$.

▶ **Example 6** (Folklore)**.** Consider the following game: Adam and Eve simultaneously and independently choose two numbers: Eve chooses $e \in \omega \cup \{\infty\}$, Adam chooses $a \in \omega \cup \{\infty\}$. Eve wins if $e < \infty$, and either $a = \infty$ or $a \leq e$.

It is easy to see that this game is not determined under mixed strategies. Intuitively, it follows from the fact that both players try to choose a finite number as big as possible.

The board $\mathtt{B}$ of the game G is depicted in Figure 3. A play $t \in \mathrm{plays}(\mathtt{B})$ consists of infinitely many independent sub-games that start in the vertices labelled by b. More precisely, the k-th sub-game starts in the node $\mathtt{L}^k\mathtt{R}$ in the tree t. Such a sub-game is split into two independent choices: Adam chooses a label, either 0 or 1, for the successor of the node labelled by a; Eve chooses a label, either 2 or 3, for the successor of the node labelled by e.

Let $a(t)$ (resp. $e(t)$) be the smallest number $k \in \omega$ such that Eve (resp. Adam) has chosen an odd label in the k-th sub-game, i.e. $\mathtt{L}^k\mathtt{RLR} \in \mathit{Dom}(t)$ (resp. $\mathtt{L}^k\mathtt{RRR} \in \mathit{Dom}(t)$). If no such number exists then $a(t)$ (resp. $e(t)$) equals ∞.

Let the winning condition L of the game G be defined as follows

$$L \stackrel{\text{def}}{=} \{t \in \text{plays}(\mathtt{B}) \mid e(t) < \infty \text{ and not } (e(t) < a(t) < \infty)\}. \tag{5}$$

It is easy to see that L is a regular language of infinite trees (to compare $a(t)$ with $e(t)$ it is enough to notice that each of these values corresponds to a node on the left-most branch of the play t). Moreover, both the conditions $e(t) < \infty$ and $e(t) < a(t) < \infty$ are open sets of plays.

Hence, the game $\mathrm{G} = \langle \mathtt{B}, L\rangle$ is a game as required in Theorem 5. Moreover, there is a clear correspondence between the pure strategies in G and the pure strategies in the game from Example 6. This correspondence extends to the mixed strategies what implies the following claim.

▶ **Claim 3.1.** *We have that* $\mathit{val}^{MA}_{\mathrm{G}} = 1$ *and* $\mathit{val}^{ME}_{\mathrm{G}} = 0$.

This concludes the proof of Theorem 5.

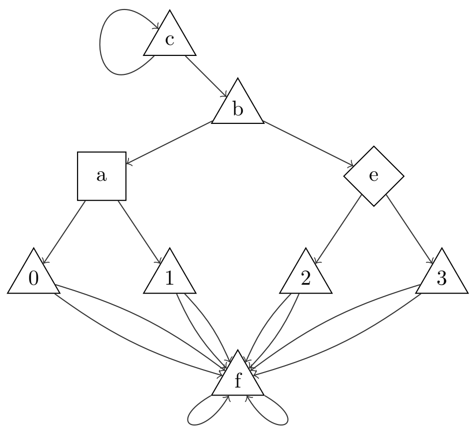

Figure 3 A branching board that is not determined under mixed strategies.

3.3 Mixed determinacy for closed sets

In this section we use Glicksberg's minimax theorem to prove that if a winning condition is a closed set of plays then the game is determined under mixed strategies.

▶ **Theorem 7.** *If* $\mathrm{G} = \langle \mathtt{B}, L\rangle$ *is a* $\{E, A, \mathcal{N}, \mathcal{B}\}$*-branching game and* L *is an arbitrary closed subset of* plays(B) *then* G *is determined under mixed strategies.*

Before we recall the statement of Glicksberg's minimax theorem, let us introduce some relevant notions. Assume that X is a metrisable topological space. We say that a function $f\colon X \to \mathbb{R}$ is *upper semi-continuous* if for every $x_0 \in X$ we have $\limsup_{x\to x_0} \leq f(x_0)$. Clearly, if $C \subseteq X$ is a closed subset of X then the characteristic function of C is upper semi-continuous. Also, a composition of a continuous function and an upper semi-continuous function is upper semi-continuous.

▶ **Theorem 8** (Glicksberg's minimax theorem [14], see also [20, pages 299–306]). *Let A, B be compact metrisable spaces and $f\colon A \times B \to \mathbb{R}$ be an upper semi-continuous function. Then the following holds*

$$\sup_{\mu} \inf_{\nu} \int_{A,B} f(a,b)\ \mu(\mathrm{d}a)\,\nu(\mathrm{d}b) = \inf_{\nu} \sup_{\mu} \int_{A,B} f(a,b)\ \mu(\mathrm{d}a)\,\nu(\mathrm{d}b), \tag{6}$$

where μ, ν range over the Borel probability measures on the sets A, B respectively.

It remains to prove that if $\mathrm{G} = \langle \mathtt{B}, L\rangle$ with $L \subseteq \text{plays}(\mathtt{B})$ closed then the function $\mathit{val}_{\mathrm{G}}\colon \Sigma^E_{\mathtt{B}} \times \Sigma^A_{\mathtt{B}} \to \mathbb{R}$ is upper semi-continuous. This function can be written as a composition of two functions. The first one maps a pair of pure strategies (σ, π) to a measure on plays(B) defined as $\mu_{(\sigma,\pi)}(T) \stackrel{\text{def}}{=} \nu^*_{\mathtt{B}}\big(\{t \in \Sigma^{\mathcal{N}}_{\mathtt{B}} \mid \sigma \cap \pi \cap t \in T\}\big)$, i.e. the $\nu^*_{\mathtt{B}}$ measure of the pre-image of the set T under the function that intersects the three strategies. This mapping is continuous, as proved by Mio in [22, Lemma 4.1.4]. The second one applies the measure $\mu_{(\sigma,\pi)}$ to the winning set $L \subseteq$ plays(B). For a closed set L this function is upper semi-continuous by [16, Corollary 17.21].

4 Computing game values

In this section we will discuss the computational complexity of determining the partial values of branching games. To be more precise, we consider the following family of problems, parametrised by the set of available positions $\mathcal{P} \subseteq \{A, E, \mathcal{N}, \mathcal{B}\}$ and the type of the value $V \in \{val^A, val^{BA}, val^{MA}, val^{ME}, val^{BE}, val^E\}$.

▶ **Problem 9** (The value V of a regular $\mathcal{P}$-branching game)**.**

- **Input** *A finitary $\mathcal{P}$-branching game* G *with the winning condition given by a non-deterministic tree automaton.*
- **Output** *Does $V > \frac{1}{2}$?*

4.1 The non-stochastic case

If no random choice is involved, i.e. the board has no *Nature*'s positions and we consider pure strategies, the values belong to the set $\{0, 1\}$ and we can compute them, as expressed by the following theorem.

▶ **Theorem 10.** *The value val^E problem of a regular $\{A, E, \mathcal{B}\}$-branching game is in 2-EXP, the value val^A problem of a regular $\{A, E, \mathcal{B}\}$-branching game is EXP-complete.*

Moreover, the value val^E problem of a regular $\{A, E, \mathcal{B}\}$-branching game is 2-EXP-complete if the winning condition is given by an alternating tree automaton.

This theorem follows from the constructions in [26], performed in a bit different language. The asymmetry in this theorem comes from the fact that in Problem 9 we assume that the winning condition of a game is given as a non-deterministic automaton. In this work we strengthen the second part of the above theorem by proving that the value val^E problem of a regular $\{A, E, \mathcal{B}\}$-branching game is *2-EXP*-hard also for non-deterministic automata. This is achieved by using the completeness result from [26] together with the following reduction. It is somehow surprising to notice that in the context of branching games one can de-alternate an automaton in polynomial time.

▶ **Theorem 11.** *There exists a polynomial time reduction that inputs a $\{A, E, \mathcal{B}\}$-branching game* G *with the winning condition given as an alternating tree automaton and constructs a $\{A, E, \mathcal{B}\}$-branching game* G′ *with the winning condition given by a non-deterministic tree automaton, such that $val^E_{\mathrm{G}} = val^E_{\mathrm{G}'}$.*

The proof is straithforward, its main idea is to split the alternation of the given automaton into two parts: the choices of Adam and the choices of Eve. In the game G′ the former choices will be done explicitly on the board while the latter choices will be performed by the non-deterministic automaton that recognises the winning condition of G′.

▶ **Corollary 12.** *val^E problem of a regular $\{A, E, \mathcal{B}\}$-branching game is 2-EXP-complete.*

Figure 4 Boards used in undecidability proofs.

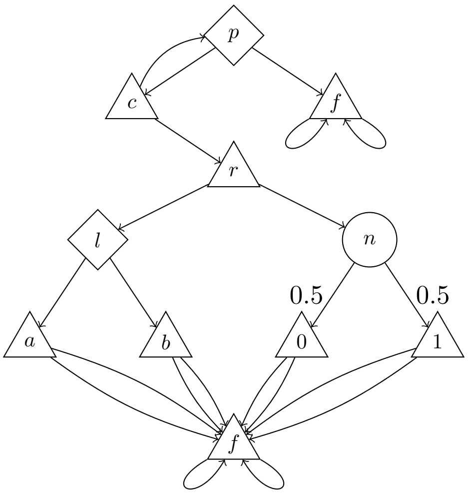

(a) A branching board used in the proof of Theorem 13.

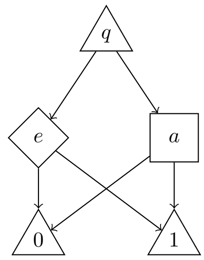

(b) A gadget used in the proof of Theorem 15. to replace *Nature*'s vertex in the board

4.2 The stochastic cases

The above decidability results hold for non-stochastic games and pure strategies. Restoring any of those features yields undecidability, as expressed by the two theorems of this section.

▶ **Theorem 13.** *For every* $V \in \{val^A, val^{BA}, val^{MA}, val^{ME}, val^{BE}, val^E\}$ *and* $P \in \{E, A\}$*, the value* V *problem of a regular* $\{P, \mathcal{N}, \mathcal{B}\}$*-branching game is undecidable.*

Observe that by Lemma 1 a $\{P, \mathcal{N}, \mathcal{B}\}$-branching game is determined under pure strategies. It means that all the six partial values are the same for such games. Thus, by the symmetry we can assume that $P = E$ and $V = val^E$.

To prove Theorem 13 we reduce the following undecidable problem, cf. [13]. It can be shown that the word problem is undecidable even if we restrict our attention to a two-letters alphabet and the so-called *very simple* automata: a non-deterministic automaton is *very simple* if from every state and every letter there are exactly two possible transitions leading to two distinct states.

▶ **Problem 14** (Word problem for VSNA)**.**

- **Input** *A very simple non-deterministic automaton* $\mathcal{A}$ *on finite words over* $\{a, b\}$*.*
- **Output** *Does there exist a finite word such that more than half of the runs of* $\mathcal{A}$ *on this word is accepting?*

We will now sketch the proof of Theorem 13. Let us take a very simple non-deterministic automaton $\mathcal{A}$ and assume that the two transitions over a letter $l \in \{a, b\}$ from a state $q \in Q^{\mathcal{A}}$ lead to the states $\delta_0(q, l)$ and $\delta_1(q, l)$. Assume that $l_0, l_1, \ldots, l_k$ is a sequence of letters $l_i \in \{a, b\}$ and $n_0, n_1, \ldots, n_k$ is a sequence of numbers $n_i \in \{0, 1\}$. These two sequences allow us to naturally define a run $\rho = \mathrm{run}(\vec{l}, \vec{n})$ of $\mathcal{A}$ over the word $l_0, \ldots, l_k$ that follows the respective transitions of $\mathcal{A}$: $\rho[0] = q_{\mathrm{I}}$ and $\rho[i+1] = \delta_{n_i}(\rho[i], l_i)$.

Consider the board $\mathtt{B}$ depicted on Figure 4a. A play on this board consists of a sequence of decisions made by Eve, whether to move from the vertex labelled l to a or to b. At every moment Eve can stop this sequence by choosing the right successor of the vertex labelled

p. For every choice of a or b by Eve, the *Nature* simultaneously chooses a number 0 or 1. Thus, a play t results in two finite or infinite sequences of the same length: $l_0, l_1, \ldots$ with $l_i \in \{a, b\}$ and $n_0, n_1, \ldots$ with $n_i \in \{0, 1\}$. Consider the following winning condition

$$L \stackrel{\text{def}}{=} \{t \in \text{plays}(\mathtt{B}) \mid \text{the sequences } \vec{l} \text{ and } \vec{n} \text{ are finite and } \text{run}(\vec{l}, \vec{n}) \text{ is accepting}\}. \quad (7)$$

Now let $\mathrm{G} = \langle \mathtt{B}, L \rangle$. It is easy to see that the winning condition L can be represented as a regular language of infinite trees. A pure strategy of Eve in G either never moves from the vertex labelled p to the vertex labelled f (in that case its value is 0) or in the opposite case it corresponds to a finite word $l_0, l_1, \ldots, l_k$. The value of such a strategy is the probability that the choices of *Nature* will represent an accepting run of $\mathcal{A}$ over the word $\vec{l}$. Thus, Eve has a pure strategy σ with $val^E_{\mathrm{G}}(\sigma) > \frac{1}{2}$ if and only if more than half of the runs of $\mathcal{A}$ over the word $\vec{l}$ produced by σ is accepting.

To complete the landscape of decidability we state.

▶ **Theorem 15.** *For every $V \in \{val^{BA}, val^{MA}, val^{ME}, val^{BE}\}$ the value V problem of a regular $\{E, A, \mathcal{B}\}$-branching game is undecidable.*

The theorem follows from the fact that the game used in the proof of Theorem 13 can be simulated on the board with *Nature*'s position replaced by the gadget depicted in Figure 4b.

5 Conclusions

In this work we have studied questions of determinacy and decidability of regular branching games. We have shown that the games are not determined even for topologically simple regular conditions. In the case of mixed determinacy, the frontier lies in the first level of the difference hierarchy of closed sets. Additionally, we have shown that the question whether the value of a given game is greater than a fixed threshold is undecidable in all non-trivial stochastic cases. In the non-stochastic cases (i.e. when the board is non-stochastic and we ask about pure strategies) we have given exact bounds on the complexity of the problem. The only remaining case is the 0-player stochastic case, i.e. the problem of computing the measure of a regular language of infinite trees.

Further work It seems interesting to understand for which classes of regular winning conditions, the branching games are determined. It was proved by the first author in [26] that the non-stochastic branching games with winning conditions given by *game automata* are determined under pure strategies. We believe that the proof can be naturally extended to the stochastic case. However, there are regular languages of infinite trees L that are not recognisable by game automata, but still all the branching games with the winning condition L are purely determined. The characterisation of such objectives poses an interesting research direction as it could give a broader class of games with decidable value problem.

On the frontier of mixed determinacy, it seems that allowing the objective to check local consistency at arbitrary depths of the tree is the cause of both the indeterminacy and the undecidability. This intuition suggests the following conjecture. We say that L is a *path language* if L is a Boolean combination of languages of the form

$$\{t \mid \text{there exists a branch of } t \text{ belonging to a regular language of infinite words } K \subseteq \Gamma^\omega\}.$$

▶ **Conjecture 16.** *If $\mathrm{G} = \langle \mathtt{B}, L \rangle$ is a branching game and L is a path language then the game* G *is determined under mixed strategies.*

The authors would like to thank Damian Niwiński for insightful discussions on the subject.

References

1 Rajeev Alur, Salvatore La Torre, and P. Madhusudan. Playing games with boxes and diamonds. In Roberto Amadio and Denis Lugiez, editors, *CONCUR 2003 - Concurrency Theory*, volume 2761 of *Lecture Notes in Computer Science*, pages 128–143. Springer Berlin Heidelberg, 2003.

2 Eugene Asarin, Julien Cervelle, Aldric Degorre, Catalin Dima, Florian Horn, and Victor Kozyakin. Entropy games. *CoRR*, abs/1506.04885, 2015. URL: `http://arxiv.org/abs/1506.04885`.

3 Robert J. Aumann. Mixed and behavior strategies in infinite extensive games. Research Memorandum 32, Economic Research Program, Princeton University, 1961.

4 Julius Richard Büchi and Lawrence H. Landweber. Solving sequential conditions by finite-state strategies. *Transactions of the American Mathematical Society*, 138:295–311, 1969.

5 Krishnendu Chatterjee and Laurent Doyen. The complexity of partial observation parity games. In *In Proc. LPAR'10, LNCS 6397*. Springer, 2010.

6 Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. A survey of partial-observation stochastic parity games. *Formal Methods in System Design*, 43(2):268–284, 2012.

7 Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, and Jean-François Raskin. Algorithms for omega-regular games with imperfect information. *CoRR*, abs/0706.2619, 2007. URL: `http://arxiv.org/abs/0706.2619`.

8 Krishnendu Chatterjee and Thomas A. Henzinger. A survey of stochastic ω-regular games. *J. Comput. Syst. Sci.*, 78(2):394–413, 2012.

9 Taolue Chen, Klaus Dräger, and Stefan Kiefer. Model checking stochastic branching processes. In *MFCS*, pages 271–282, 2012.

10 Thomas Colcombet and Damian Niwiński. On the positional determinacy of edge-labeled games. *Theor. Comput. Sci.*, 352(1–3):190–196, 2006.

11 Stefan Dziembowski, Marcin Jurdziński, and Igor Walukiewicz. How much memory is needed to win infinite games? In *LICS*, pages 99–110, 1997.

12 Allen Emerson and Charanjit Jutla. Tree automata, mu-calculus and determinacy. In *FOCS'91*, pages 368–377, 1991.

13 Hugo Gimbert and Youssouf Oualhadj. Probabilistic automata on finite words: Decidable and undecidable problems. In *ICALP (2)*, pages 527–538, 2010.

14 Irving Leonard Glicksberg. *Minimax Theorem for Upper and Lower Semi-continuous Payoffs*. Memorandum (Rand Corporation). Rand Corporation, 1950.

15 Tomasz Gogacz, Henryk Michalewski, Matteo Mio, and Michał Skrzypczak. Measure properties of game tree languages. In *MFCS*, pages 303–314, 2014.

16 Alexander Kechris. *Classical descriptive set theory*. Springer-Verlag, New York, 1995.

17 Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.

18 Donald A. Martin. The determinacy of Blackwell games. *The Journal of Symbolic Logic*, 63(4):1565–1581, 1998.

19 Robert McNaughton. Infinite games played on finite graphs. *Annals of Pure and Applied Logic*, 65(2):149–184, 1993.

20 P. Wolfe Melvin Dresher, A. W. Tucker. *Contributions to the Theory of Games*. Number 3 in Annals of mathematics studies. Princeton University Press, 1957.

21 Henryk Michalewski and Matteo Mio. On the problem of computing the probability of regular sets of trees. In *FSTTCS*, pages 489–502, 2015.

22 Matteo Mio. *Game Semantics for Probabilistic μ-Calculi*. PhD thesis, University of Edinburgh, 2012.

23 Matteo Mio. On the equivalence of game and denotational semantics for the probabilistic mu-calculus. *Logical Methods in Computer Science*, 8(2), 2012.

24 Andrzej W. Mostowski. Games with forbidden positions. Technical report, University of Gdańsk, 1991.

25 A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proceedings of the 16th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '89, pages 179–190, New York, NY, USA, 1989. ACM. doi:10.1145/75277.75293.

26 Marcin Przybyłko. Tree games with regular objectives. In *GandALF*, pages 231–244, 2014.

27 Michael Oser Rabin. Decidability of second-order theories and automata on infinite trees. *Trans. of the American Math. Soc.*, 141:1–35, 1969.

28 Wolfgang Thomas. Languages, automata, and logic. In *Handbook of Formal Languages*, pages 389–455. Springer, 1996.

29 Wolfgang Thomas and Helmut Lescow. Logical specifications of infinite computations. In *REX School/Symposium*, pages 583–621, 1993.

30 Moshe Y. Vardi. An automata-theoretic approach to linear temporal logic. In Faron Moller and Graham Birtwistle, editors, *Logics for Concurrency*, volume 1043 of *Lecture Notes in Computer Science*, pages 238–266. Springer Berlin Heidelberg, 1996.

A Defining the measure $\eta^*_{\mathtt{B}}$

In this section we provide a precise definition of the measure $\eta^*_{\mathtt{B}}$ induced by a board $\mathtt{B}$. The aim of this measure is to guarantee the following equation. Consider a direction $d \in \{\mathtt{L}, \mathtt{R}\}$ and a node $u \in \{\mathtt{L}, \mathtt{R}\}^*$ such that u corresponds to a position of $\mathcal{N}$ in $\mathtt{B}$, i.e. $s_u(v_{\mathrm{I}}) = v \in V$ and $\rho(v) = \mathcal{N}$.

$$\mathbb{P}_{t \in \Sigma^{\mathcal{N}}_{\mathtt{B}}}\big(ud \in t \mid u \in t\big) = \eta(v)(d). \tag{8}$$

In other words, if a play reaches a vertex $v \in V$ that belongs to $\mathcal{N}$ then the probability that a random strategy of $\mathcal{N}$ will choose a direction $d \in \{\mathtt{L}, \mathtt{R}\}$ equals $\eta(v)(d)$.

A simple definition of the measure $\eta^*_{\mathtt{B}}$ can be given by Caratheodori's theorem. Here we take a different approach. Consider the space X of all functions $x\colon \{\mathtt{L}, \mathtt{R}\}^* \to [0, 1]$, where $[0, 1] \subseteq \mathbb{R}$ is the closed interval. There is a natural probability measure μ_X on X that is the countable independent product of the Lebesgue measures on $[0, 1]$. Let us fix a board $\mathtt{B}$ and define $f_{\mathtt{B}}\colon X \to \Sigma^{\mathcal{N}}_{\mathtt{B}}$ inductively: for $u \in \{\mathtt{L}, \mathtt{R}\}^*$ and $v_u = s_u(v_{\mathrm{I}})$:

- $f_{\mathtt{B}}(x)(\varepsilon) = \lambda(v_\varepsilon)$,
- if $u \in \mathit{Dom}(f_{\mathtt{B}}(x))$ and $v_u \notin V_{\mathcal{N}}$ then for $d = \mathtt{L}, \mathtt{R}$ we put $f_{\mathtt{B}}(x)(ud) = \lambda(v_{ud})$,
- if $u \in \mathit{Dom}(f_{\mathtt{B}}(x))$, $v_u \in V_{\mathcal{N}}$, and $x(u) < \eta(v_u)(\mathtt{L})$ then we put $f_{\mathtt{B}}(x)(u\mathtt{L}) = \lambda(v_{u\mathtt{L}})$ and let $u\mathtt{R} \notin \mathit{Dom}(f_{\mathtt{B}}(x))$,
- if $u \in \mathit{Dom}(f_{\mathtt{B}}(x))$, $v_u \in V_{\mathcal{N}}$, and $x(u) \geq \eta(v_u)(\mathtt{L})$ then we let $u\mathtt{L} \notin \mathit{Dom}(f_{\mathtt{B}}(x))$ and we put $f_{\mathtt{B}}(x)(u\mathtt{R}) = \lambda(v_{u\mathtt{R}})$.

In other words, the function $f_{\mathtt{B}}$ resolves the choices of the *Nature* depending on the randomly chosen values of x in the respective nodes of the tree. It is easy to see that the function $f_{\mathtt{B}}$ is measurable (in fact Borel). Now let us define $\eta^*_{\mathtt{B}}(A)$ for a measurable set $A \subseteq \Sigma^{\mathcal{N}}_{\mathtt{B}}$ as

$$\eta^*_{\mathtt{B}}(A) = \mu_X\big(f_{\mathtt{B}}^{-1}(A)\big).$$

It is easy to verify that $\eta^*_{\mathtt{B}}$ defined this way satisfies (8).

B Proof of Lemma 1

In this section we prove Lemma 1.

▶ **Lemma 1.** *Let* G *be a branching game. If σ_m is Eve's mixed strategy then*

$$\inf_{\pi_m \in \Sigma^{MA}_{\mathtt{B}}} \mathit{val}_{\mathrm{G}}(\sigma_m, \pi_m) = \inf_{\pi_b \in \Sigma^{BA}_{\mathtt{B}}} \mathit{val}_{\mathrm{G}}(\sigma_m, \pi_b) = \inf_{\pi \in \Sigma^{A}_{\mathtt{B}}} \mathit{val}_{\mathrm{G}}(\sigma_m, \pi)$$

Dually, the same holds for mixed strategies of Adam if we replace inf *with* sup *and A with E.*

By duality we focus only on the case of Eve. Let us fix a mixed strategy σ_m of Eve. Clearly the $\leq$-inequalities hold because of the containment of the sets of strategies of Adam. Assume that $\inf_{\pi_m \in \Sigma^{MA}_{\mathtt{B}}} \mathit{val}_{\mathrm{G}}(\sigma_m, \pi_m) < \inf_{\pi \in \Sigma^{A}_{\mathtt{B}}} \mathit{val}_{\mathrm{G}}(\sigma_m, \pi)$. It means that there exists a single mixed strategy of Adam $\pi_m \in \Sigma^{MA}_{\mathtt{B}}$ such that $\mathit{val}_{\mathrm{G}}(\sigma_m, \pi_m) < \inf_{\pi \in \Sigma^{A}_{\mathtt{B}}} \mathit{val}_{\mathrm{G}}(\sigma_m, \pi)$. Let $\pi \in \Sigma^{A}_{\mathtt{B}}$ be a pure strategy of Adam and let us define

$$f(\pi) \stackrel{\text{def}}{=} \mathit{val}_{\mathrm{G}}(\sigma_m, \pi) = \int_{\Sigma^{E}_{\mathtt{B}}, \Sigma^{\mathcal{N}}_{\mathtt{B}}} \Phi\big(\mathrm{eval}_{\mathrm{G}}(\sigma, \pi, \eta)\big)\ \sigma_m(\mathrm{d}\sigma)\, \eta^*_{\mathtt{B}}(\mathrm{d}\eta).$$

By Fubini's theorem, we know that

$$\mathit{val}_{\mathrm{G}}(\sigma_m, \pi_m) = \int_{\Sigma^{A}_{\mathtt{B}}} f(\pi)\ \pi_m(\mathrm{d}\pi).$$

It means that under our assumption we have $\int_{\Sigma_{\mathrm{B}}^{A}} f(\pi)\ \pi_m(\mathrm{d}\pi) < \inf_{\pi\in\Sigma_{\mathrm{B}}^{A}} f(\pi)$. This is a contradiction, since for every probability measure μ on a Polish topological space X and a μ-measurable function $f\colon X \to \mathbb{R}^+$ we have

$$\int_X f(x)\ \mu(\mathrm{d}x) \geq \left(\inf_{x\in X} f(x)\right)\cdot \mu(X) = \inf_{x\in X} f(x).$$

C Behavioural indeterminacy

In this section we compute the behavioural values of the game from Subsection 3.1 stated in Example 4. Fix a behavioural strategy σ_b of Eve that is characterised by two values $p_1, p_2 \in [0,1]$. Notice that there are four pure strategies of Adam, indexed by $i, j \in \{0,1\}$: the strategy $\pi_{i,j}$ satisfies $x_3(\pi_{i,j}) = i$ and $x_4(\pi_{i,j}) = j$. Their values are:

- $val_{\mathrm{G}}(\sigma_b, \pi_{00}) = p_1 \cdot p_2$;
- $val_{\mathrm{G}}(\sigma_b, \pi_{01}) = 1$;
- $val_{\mathrm{G}}(\sigma_b, \pi_{10}) = 1$;
- $val_{\mathrm{G}}(\sigma_b, \pi_{11}) = (1-p_1)\cdot(1-p_2)$.

Since Eve wants to maximise the value we infer that

$$val_{\mathrm{G}}^{BE} = \max_{p_1,p_2\in[0,1]} \min\big(p_1p_2,\ (1-p_1)(1-p_2)\big). \tag{9}$$

We claim that the value (9) is $\frac{1}{4}$.

First of all, we notice that for $p_1 = p_2 = \frac{1}{2}$ we have $(1-p_1)\cdot(1-p_2) = p_1\cdot p_2 = \frac{1}{4}$. Hence, $val_{\mathrm{G}}^{BE} \geq \frac{1}{4}$.

Now, if $p_2 \leq 1-p_1$ then $val_{\mathrm{G}}^{BE} = p_1\cdot p_2 \leq p_1\cdot(1-p_1) \leq \frac{1}{4}$. Similarly, if $p_2 \geq 1-p_1$ then $val_{\mathrm{G}}^{BE} = (1-p_1)\cdot(1-p_2) \leq p_2\cdot(1-p_2) \leq \frac{1}{4}$. Thus, $val_{\mathrm{G}}^{BE} \leq \frac{1}{4}$ and finally $val_{\mathrm{G}}^{BE} = \frac{1}{4}$.

A similar analysis infers that

$$\begin{aligned}
&val_{\mathrm{G}}^{BA} = \\
&= \min_{p_1,p_2\in[0,1]} \max\big(p_1p_2 + p_1(1{-}p_2) + p_2(1{-}p_1),\ (1{-}p_1)(1{-}p_2) + p_1(1{-}p_2) + p_2(1{-}p_1)\big)\\
&= \min_{p_1,p_2\in[0,1]} \max\big(1-(1{-}p_1)(1{-}p_2),\ 1-p_1p_2\big)\\
&= 1 - \max_{p_1,p_2\in[0,1]} \min\big((1{-}p_1)(1{-}p_2),\ p_1p_2\big)\\
&= 1-\frac{1}{4} = \frac{3}{4}.
\end{aligned}$$

D Mixed indeterminacy

In this section we prove Claim 3.1 from Subsection 3.2.

▶ **Claim 3.1.** *We have that* $val_{\mathrm{G}}^{MA} = 1$ *and* $val_{\mathrm{G}}^{ME} = 0$.

We will provide an argument that $val_{\mathrm{G}}^{ME} = 0$, the argument for Adam is entirely dual. Consider any mixed strategy σ_m of Eve in the game G and take $\epsilon > 0$. We will prove that $val_{\mathrm{G}}^{E}(\sigma_m) < \epsilon$ which will imply that $val_{\mathrm{G}}^{E}(\sigma_m) = 0$.

Notice that for every pure strategy σ of Eve the value $e(\sigma)$ as defined in Subsection 3.2 is well-defined (it depends only on the pure strategy of Eve). The same holds for the pure strategies π of Adam and the value $a(\pi)$.

For every $i \in \omega \cup \{\infty\}$ let us define

$$P_i \stackrel{\text{def}}{=} \sigma_m\big(\{\sigma \in \Sigma^E_{\mathtt{B}} \mid e(\sigma) = i\}\big),$$

i.e. the probability that the mixed strategy σ_m will choose a pure strategy σ satisfying $e(\sigma) = i$. Clearly $0 \leq P_i \leq 1$ and $\sum_{i \in \omega \cup \{\infty\}} P_i = 1$. Let $I \in \omega$ be such a number that

$$\sum_{i \in \omega, i \geq I} P_i < \epsilon.$$

Such a number I exists because the sum $\sum_{i \in \omega} P_i$ is finite.

Consider the pure strategy π of Adam that satisfies $a(\pi) = I$, i.e. it chooses the label 1 only in the I-th sub-game. By the definition of the winning condition of G we obtain that:

$$\mathit{val}_{\mathrm{G}}(\sigma_m) \leq \mathit{val}_{\mathrm{G}}(\sigma_m, \pi) \leq \sum_{i \in \omega, i \geq I} P_i < \epsilon.$$

E Adam's value problem for non-stochastic boards

We claim that the following theorem is a consequence of the techniques from [26].

▶ **Theorem 17.** *(a) The value val^E problem of a regular $\{A, E, \mathcal{B}\}$-branching game is 2-EXP-complete if the winning condition is given by an alternating tree automaton.*
(b) Moreover, the value val^A problem of a regular $\{A, E, \mathcal{B}\}$-branching game is EXP-complete.

For completeness we will describe how to recover those results. The following are the two crucial lemmas from [26] expressed in our notation.

▶ **Lemma 18** ([26, Lemma 5])**.** *For every finitary $\{A, E, \mathcal{B}\}$-branching game* G *with a regular winning condition* $\mathrm{L}(\mathcal{A})$*, the set of winning strategies of Eve is regular. Moreover, it can be recognised by an alternating finite automaton of size exponential in the size of $\mathcal{A}$ and polynomial in the size of G, where $\mathcal{A}$ is an alternating automaton recognising the winning condition* $\mathrm{L}(\mathcal{A})$*.*

▶ **Lemma 19** ([26, Lemma 3])**.** *The value val^E problem of a regular $\{A, E, \mathcal{B}\}$-branching game is EXP-hard.*

The reduction in the proof of [26, Lemma 3] uses $\{A, \mathcal{B}\}$-branching games which are determined ($\mathit{val}^E = \mathit{val}^A$) and, thus, *EXP*-hardness in *(b)* follows.

The hardness in *(a)* is recovered from the proof [26, Corollary 1]. Solving standard two player games with objectives defined by an *LTL* formulae is *2-EXP*-complete (cf. e.g. [25, 1]). Since an *LTL* formula can be translated into an alternating finite automaton on infinite words in polynomial time (e.g., see [30]) and standard two player games are a special case of branching games; the lower bound of *(a)* immediately follows.

The doubly-exponential-time algorithm in part *(a)* is a simple application of [26, Lemma 5], recovering the exponential-time algorithm in part *(b)* is more involved.

The doubly-exponential-time algorithm that can be recovered from [26, Lemma 5] checks the non-emptiness of an alternating automaton recognising the set of Eve's winning strategies and can be described in the following steps.

Given finitary board $\mathtt{B}$ and an alternating tree automaton $\mathcal{A}$ do

1. compute a non-deterministic automaton $\mathcal{A}_2$ recognising the complement of $\mathrm{L}(\mathcal{A})$,

2. compute a non-deterministic automaton $\mathcal{A}_3$ recognising the set of Eve's strategies that are not winning in $\langle \mathtt{B}, \mathrm{L}(\mathcal{A})\rangle$,
3. compute an automaton $\mathcal{A}_4$ recognising the complement of the language $\mathrm{L}(\mathcal{A}_3)$,
4. check the non-emptiness of $\mathrm{L}(\mathcal{A}_4)$.

Only Step 1. of this procedure requires an exponential blow-up in the size of the automaton, the remaining automata can be constructed in polynomial time. Therefore, the size of the automaton in Step 4. is exponential in terms of the size of the original automaton. Since the non-emptiness of an alternating finite automaton on trees is *EXP*-complete, we obtain the *2-EXP* upper bound. Still, notice that if instead of the automaton recognising the winning set the algorithm receives a non-deterministic automaton recognising the complement of the wining set as the input, we can avoid the blow-up. This is exactly the case when we compute val^A of a regular $\{A, E, \mathcal{B}\}$-branching game and, thus, we infer the *EXP* upper bound.

F Hardness

In this section we prove Theorem 11 from Subsection 4.1.

▶ **Theorem 11.** *There exists a polynomial time reduction that inputs a $\{A, E, \mathcal{B}\}$-branching game* G *with the winning condition given as an alternating tree automaton and constructs a $\{A, E, \mathcal{B}\}$-branching game* G' *with the winning condition given by a non-deterministic tree automaton, such that $val^E_{\mathrm{G}} = val^E_{\mathrm{G}'}$.*

F.1 Automata

Before we prove the theorem, we will introduce formally the notions of a non-deterministic and alternating automata over infinite trees.

Parity games Parity games [24, 12] are one of the crucial concepts in the theory of regular languages of infinite trees. For the sake of succinctness we will formalise them as a very special case of the branching games. A *parity game* is a $\{E, A\}$-branching game that is additionally equipped with a function $\alpha\colon V \to \{i, i+1, \dots, j\} \subseteq \omega$ that assigns *priorities* to vertices. A play t of a parity game has the shape of a unique infinite branch that corresponds to a sequence of vertices $v_{\mathrm{I}} = v_0, v_1, \dots$. Such a play is won by Eve (i.e. has value 1) if the limes inferior of the values $\alpha(v_i)$ is even. Otherwise $\Phi(t) = 0$. The notions of pure strategies for the respective players are the same as in the case of general branching games. We call a pure strategy τ of $P \in \{E, A\}$ *positional* if the decision made by P depends only on the current vertex $v \in V$. In other words, for every two $u, w \in Dom(\tau)$ such that $s_u(v_{\mathrm{I}}) = s_w(v_{\mathrm{I}}) \in V_P$, either $u\mathtt{L}, w\mathtt{L} \in Dom(\tau)$ or $u\mathtt{R}, w\mathtt{R} \in Dom(\tau)$. Thus, a positional strategy of P can be represented as a function $\tau'\colon V_P \to \{\mathtt{L}, \mathtt{R}\}$. A classical result [24, 12] states that for every parity game one of the players has a pure positional winning strategy.

Automata An *Alternating Tree Automaton* is a tuple $\mathcal{A} = \langle Q, \Gamma, \delta, \alpha, q_{\mathrm{I}}\rangle$ where Q is a finite set of *states*; Γ is a finite alphabet; α is a function assigning *priority* $\alpha(q) \in \{i, i+1, \dots, j\}$ to every state $q \in Q$; q_{I} is the *initial state*; and δ is a *transition function* mapping pairs $\langle q, a\rangle \in Q \times \big(\Gamma \cup \{\star\}\big)$ to formulae ϕ built using the grammar

$$\phi ::= \phi \wedge \phi \mid \phi \vee \phi \mid (p, d) \mid \top \mid \bot,$$

where $d \in \{\mathtt{L}, \mathtt{R}\}$, $p \in Q$, and $\star$ is a fresh letter, $\star \notin \Gamma$.

We say that a formula φ is

- atomic if φ is of the form (p,d), $\top$, or $\bot$; where $p \in Q$, $d \in \{\mathtt{L},\mathtt{R}\}$,
- an $\wedge$-*formula* if φ is atomic or of the form $(p,\mathtt{L}) \wedge (q,\mathtt{R})$ where $p,q \in Q$,
- an $\vee$-*formula* if φ is atomic or of the form $(p,\mathtt{L}) \vee (q,\mathtt{R})$ where $p,q \in Q$.

An automaton $\mathcal{A}$ is called
- *deterministic* if the formulae are $\wedge-$formulae,
- *game* if the formulae are either $\vee-$ or $\wedge-$formulae,
- *non-deterministic* if the formulae are disjunctions of $\wedge-$formulae.

Let t be a tree over an alphabet Γ, $\star \notin \Gamma$ be a fresh letter, and $t' \colon \{\mathtt{L},\mathtt{R}\}^* \to \Gamma \cup \{\star\}$ be an extension of t, i.e. $t \subseteq t'$ and for $u \in \{\mathtt{L},\mathtt{R}\}^* \setminus \mathit{Dom}(t)$ we have that $t'(u) = \star$. Assume that Δ is the set of all sub-formulae of transitions of an alternating tree automaton $\mathcal{A}$. $\mathcal{A}$ *accepts* the tree t if Eve has a winning strategy in the parity game $\mathrm{G}(\mathcal{A},t) = \langle V, s_{\mathtt{L}}, s_{\mathtt{R}}, \rho, \alpha, v_{\mathrm{I}}\rangle$ defined as:
- $V \stackrel{\text{def}}{=} \Delta \times \{\mathtt{L},\mathtt{R}\}^*$,
- $v_{\mathrm{I}} \stackrel{\text{def}}{=} \langle \delta(q_{\mathrm{I}}, t(\varepsilon)), \varepsilon\rangle$.
- If $m = \max_{q\in Q} \alpha(q)$, then ρ, α, $s_{\mathtt{L}}$, and $s_{\mathtt{L}}$ are defined as follows: for each $v = \langle \psi, w\rangle \in V$
  - if $\psi = \psi_{\mathtt{L}} \vee \psi_{\mathtt{R}}$, then $\rho(v) = E$, $s_d(v) = \langle \psi_d, w\rangle$, and $\alpha(v) = m$, for $d \in \{\mathtt{L},\mathtt{R}\}$;
  - if $\psi = \psi_{\mathtt{L}} \wedge \psi_{\mathtt{R}}$, then $\rho(v) = A$, $s_d(v) = \langle \psi_d, w\rangle$, and $\alpha(v) = m$, for $d \in \{\mathtt{L},\mathtt{R}\}$;
  - if $\psi = (d,q)$ then $\rho(v) = E$, $s_{\mathtt{L}}(v) = s_{\mathtt{R}}(v) = \big(\delta(q, t'(w)), wd\big)$, and $\alpha(\psi,w) = \alpha(q)$;
  - if $\psi = \top$ then $\rho(v) = E$, $s_{\mathtt{L}}(v) = s_{\mathtt{R}}(v) = v$, and $\alpha(\psi,w) = 0$;
  - if $\psi = \bot$ then $\rho(v) = E$, $s_{\mathtt{L}}(v) = s_{\mathtt{R}}(v) = v$, and $\alpha(\psi,w) = 1$.

We denote the set of trees accepted by $\mathcal{A}$ as $\mathrm{L}(\mathcal{A})$ and call the *language recognised* by $\mathcal{A}$.

F.2 Construction of the board $\mathtt{B}'$

Consider a $\{A,E,\mathcal{B}\}$-branching game $\mathrm{G} = \langle \mathtt{B}, \mathrm{L}(\mathcal{A})\rangle$ where the winning condition $\mathrm{L}(\mathcal{A})$ is given by an alternating tree automaton $\mathcal{A}$. The val^E problem for G asks if

$$\exists\sigma\,\forall\pi.\ \sigma\cap\pi \in \mathrm{L}(\mathcal{A}).$$

Since the automaton $\mathcal{A}$ is alternating, the question whether $\sigma\cap\pi \in \mathrm{L}(\mathcal{A})$ can be written equivalently as:
- $\exists\bar\sigma\,\forall\bar\pi.\ \bar\sigma$ wins with $\bar\pi$ in $\mathrm{G}(\mathcal{A},\sigma\cap\pi)$,
- $\forall\bar\pi\,\exists\bar\sigma.\ \bar\sigma$ wins with $\bar\pi$ in $\mathrm{G}(\mathcal{A},\sigma\cap\pi)$,

where $\bar\sigma$ and $\bar\pi$ range over pure strategies of the respective players in the game $\mathrm{G}(\mathcal{A},\sigma\cap\pi)$. Additionally, since parity games are positionally determined, in both cases we can restrict to positional pure strategies $\bar\sigma$ and $\bar\pi$. By incorporating the latter condition, we need to solve the following question:

$$\exists\sigma\,\forall\pi\,\forall\bar\pi\,\exists\bar\sigma.\ \bar\sigma \text{ wins with } \bar\pi \text{ in } \mathrm{G}(\mathcal{A},\sigma\cap\pi).$$

In this section we will construct another board $\mathtt{B}'$ and a non-deterministic tree automaton $\mathcal{A}'$ such that a pure strategy of Adam over $\mathtt{B}'$ will encode both, a pure strategy π over $\mathtt{B}$ and a positional strategy $\bar\pi$ in $\mathrm{G}(\mathcal{A},\sigma\cap\pi)$. The non-deterministic automaton $\mathcal{A}$ will guess a positional strategy $\bar\sigma$ of Eve in $\mathrm{G}(\mathcal{A},\sigma\cap\pi)$. Thus, the val^E problem for G' will be equivalent to asking whether:

$$\exists\sigma\,\forall(\pi,\bar\pi).\ \sigma\cap\pi\cap\bar\pi \in \mathrm{L}(\mathcal{A}').$$

Assume that the set of states of $\mathcal{A}$ is $Q = \{q_0, q_1, \ldots, q_n\}$, its initial state is q_{I} and Δ is the set of all sub-formulae in the transitions of $\mathcal{A}$. The alphabet Γ' of the new board $\mathtt{B}'$

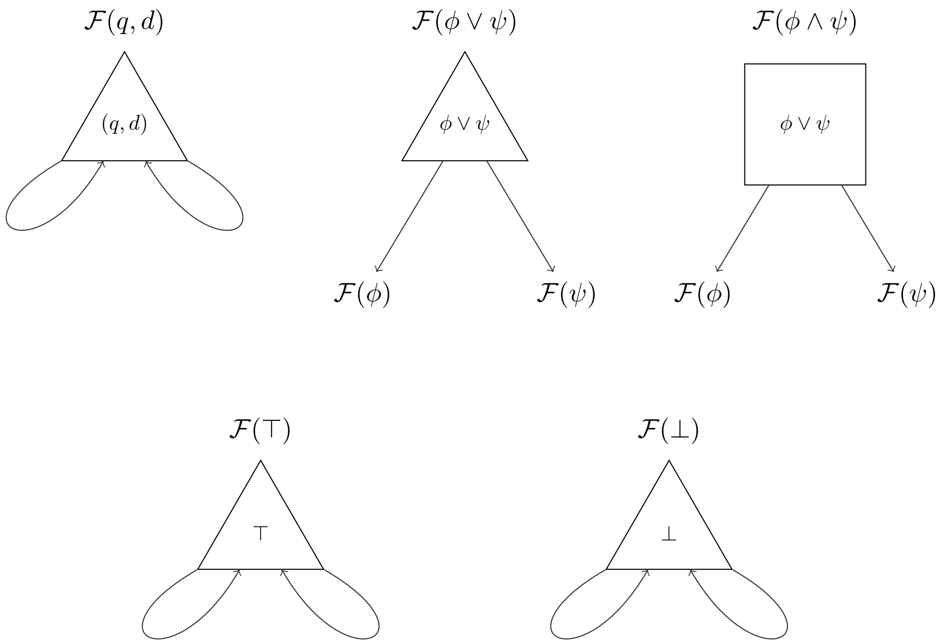

Figure 5 The inductive construction of the sub-board $\mathcal{F}(\delta)$ for $\delta \in \Delta$.

will be the disjoint union of the original alphabet Γ, the set $Q \times \Gamma$, the set Δ, and a special symbol f. Our aim is to replace each vertex $v \in V$ on the board $\mathtt{B}$ with $\lambda(v) = a$ by a gadget $\mathcal{G}(v)$ that simulates all the possible transitions of $\mathcal{A}$ over a.

First let us define inductively $\mathcal{F}(\delta)$ for $\delta \in \Delta$ as depicted in Figure 5—the atomic transitions lead to trivial sub-games looping in a vertex labelled by f, the disjunctions are replaced by branching vertices, and the conjunctions are replaced by Adam's vertices.

Now, $\mathtt{B}'$ is obtained from $\mathtt{B}$ by performing the replacement depicted in Figure 6—each vertex v of $\mathtt{B}$ that was labelled by a is replaced by a gadget listing all the states of $\mathcal{A}$ with all the transitions from these states over the letter a.

F.3 Construction of the automaton $\mathcal{A}'$

We will now describe the automaton $\mathcal{A}'$. Its aim is to simulate a play of $\mathrm{G}(\mathcal{A}, t)$ over the extended arena $\mathtt{B}'$. The only essential non-determinism of $\mathcal{A}'$ will be available when a letter of the form $\phi \vee \psi$ is read. In that case the automaton will guess a choice of Eve (from the strategy $\bar{\sigma}$) and follow the respective sub-tree generated by the sub-boards $\mathcal{F}(\phi)$ or $\mathcal{F}(\psi)$. After resolving the current transition $\delta(q, a)$ (represented as $\mathcal{F}(\delta(q, a))$) when the automaton finally reaches a vertex labelled by an atomic formula ψ (e.g. (q, d)) it needs to pass this knowledge to the successive node denoted by the hexagon in Figure 6. This will be achieved by guessing in advance that the resolved atomic formula will be ψ.

Let the set of states of $\mathcal{A}'$ be $Q \cup Q \times \{\mathtt{L}, \mathtt{R}\} \cup \{\top, \bot\}$. The states from Q will denote the fact that a transition of $\mathcal{A}$ has not been resolved yet, while the states from $Q \times \{\mathtt{L}, \mathtt{R}\} \cup \{\top, \bot\}$ will denote the fact that we already have played a finite game over the formulae in Δ. The initial state of $\mathcal{A}'$ is q_{I}. Let $m = \max_{q \in Q} \alpha(q)$ and let $\alpha'(q, d) = \alpha'(\top) = \alpha'(\bot) = m$. For $q \in Q$ let $\alpha'(q) = \alpha(q)$.

Before we list the transitions of $\mathcal{A}'$ we describe the operation of $\mathcal{A}'$ informally:

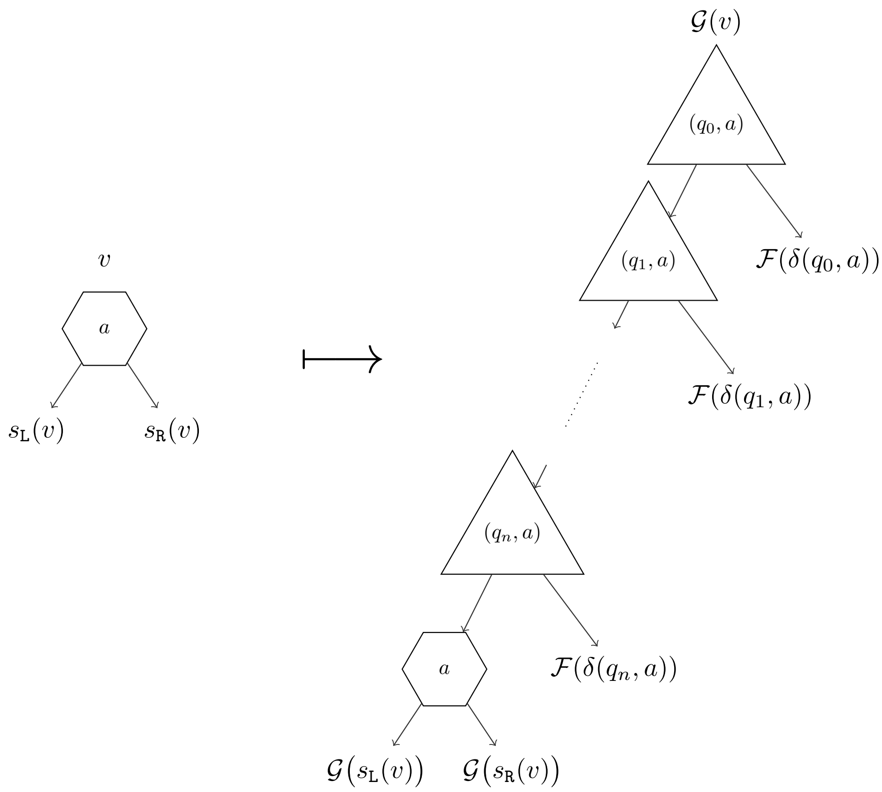

Figure 6 The transformation on the board $\mathtt{B}$ to obtain $\mathtt{B}'$. The hexagon represents an arbitrary vertex v of $\mathtt{B}$ that is labelled by $a \in \Gamma$. This vertex is replaced by a gadget $\mathcal{G}(v)$ in $\mathtt{B}'$, as depicted on the right-hand side of the figure. The gadget lists all the states $q_0, \ldots, q_n$ of $\mathcal{A}$ and for each state contains the sub-board $\mathcal{F}\big(\delta(q_i, a)\big)$ that represents the formula of the transition of $\mathcal{A}$ over a from q. Finally, the left-most branch of the gadget $\mathcal{G}(v)$ reaches a copy of the vertex v that leads to the gadgets corresponding to the successors of v in $\mathtt{B}$.

- It enters a component $\mathcal{G}(v)$ as depicted in Figure 6 in a state q.
- It passes along the left-most branch of $\mathcal{G}(v)$ until a letter (q, a) is reached.
- It guesses the atomic formula that will be reached when resolving the transition $\delta(q, a)$ (i.e. the sub-board $\mathcal{F}(\delta(q, a))$).
- It passes the guessed atomic formula along the left-most branch of $\mathcal{G}(v)$ until reaching the final node labelled by a letter $a \in \Gamma$.
- At the same time it passes the guessed atomic formula down the subtree of $\mathcal{F}(\delta(q, a))$.
- Since the conjunctive formulae of the form $\phi \wedge \psi$ were translated into Adam's positions, they are already resolved—i.e. the nodes of the tree labelled by them are uniquely branching.
- The disjunctive formulae of the form $\phi \vee \psi$ are not resolved yet and the automaton resolves them using non-determinism.
- When it reaches a node labelled by an atomic formula, it checks that the formula is the same as its state.

The automaton will have the following transitions:

- From a state q over a letter (q', a) with $q' \neq q$ it will deterministically take the transition $(q, \mathtt{L})$, following the left-most path in the gadget depicted in Figure 6.
- From a state q over a letter (q, a) it will non-deterministically take one of the transitions $\big(\theta, \mathtt{L}\big) \wedge \big(\theta, \mathtt{R}\big)$ for $\theta \in Q \times \{\mathtt{L}, \mathtt{R}\} \cup \{\top, \bot\}$. Such a transition corresponds to guessing that the formula $\delta(q, a)$ will be resolved to an atomic formula θ) (the $\big(\theta, \mathtt{R}\big)$ part) and continuing the computation on the left-most path in the state θ.
- From a state $\theta \in Q \times \{\mathtt{L}, \mathtt{R}\} \cup \{\top, \bot\}$ over a letter (q', a') it will deterministically take the transition $\big(\theta, \mathtt{L}\big)$, following the left-most path in the gadget depicted in Figure 6.
- From a state $\theta \in Q \times \{\mathtt{L}, \mathtt{R}\} \cup \{\top, \bot\}$ over a letter $\phi \vee \psi$ it will non-deterministically take one of the transitions $\big(\theta, d\big)$ for $d \in \{\mathtt{L}, \mathtt{R}\}$.
- From a state $\theta \in Q \times \{\mathtt{L}, \mathtt{R}\} \cup \{\top, \bot\}$ over a letter $\phi \wedge \psi$ it will non-deterministically take one of the transitions $\big(\theta, d\big)$ for $d \in \{\mathtt{L}, \mathtt{R}\}$—the node labelled $\phi \wedge \psi$ has exactly one child in a play and we need to follow that successor.
- From a state $\theta \in Q \times \{\mathtt{L}, \mathtt{R}\} \cup \{\top, \bot\}$ over a letter $\theta' \in Q \times \{\mathtt{L}, \mathtt{R}\} \cup \{\top, \bot\}$ with $\theta \neq \theta'$ the automaton takes the transition $\bot$.
- From a state $\theta \in Q \times \{\mathtt{L}, \mathtt{R}\} \cup \{\top, \bot\}$ over the letter θ the automaton takes the transition $\top$.
- From a state $(q, d) \in Q \times \{\mathtt{L}, \mathtt{R}\}$ over a letter from the original alphabet $a \in \Gamma$ the automaton deterministically takes the transition (q, d), moving in the direction d to the state $q \in Q$.
- From a state $\theta \in \{\top, \bot\}$ over a letter from the original alphabet $a \in \Gamma$ the automaton deterministically takes the transition θ.
- For states and letters not listed above, the automaton takes the rejecting transition.

F.4 Correctness of the construction

Let $\mathrm{G}' = \langle \mathtt{B}', \mathrm{L}(\mathcal{A}') \rangle$. In this section we prove the correctness of the construction of G', as expressed by the following claim.

▶ **Claim F.1.** *Eve has a winning strategy in the original branching game* G *if and only if she has a winning strategy in the game* G'.

First assume that σ is a pure winning strategy of Eve in the original game G. Since Eve has no additional choices over the board $\mathtt{B}'$, the strategy σ can be naturally interpreted as a pure strategy σ' over the board $\mathtt{B}'$. Consider a pure strategy π' of Adam over the board $\mathtt{B}'$. We will prove that $\sigma' \cap \pi' \in \mathrm{L}(\mathcal{A}')$. Notice that the strategy π' consists of two parts:

- it encodes a pure strategy π of Adam over the board $\mathtt{B}$,
- the choices made by π' in the vertices labelled by the elements of Δ encode a positional strategy $\bar{\pi}$ of Adam in the game $\mathrm{G}(\mathcal{A}, \sigma \cap \pi)$.

Since the strategy σ is winning, we know that $\sigma \cap \pi \in \mathrm{L}(\mathcal{A})$. Therefore, there exists a positional strategy $\bar{\sigma}$ of Eve that wins with $\bar{\pi}$ in the game $\mathrm{G}(\mathcal{A}, \sigma \cap \pi)$. We can use $\bar{\sigma}$ to define an accepting run of $\mathcal{A}'$ over the tree $\sigma' \cap \pi'$. Therefore, we have proven that $\sigma' \cap \pi' \in \mathrm{L}(\mathcal{A}')$.

Now consider a pure winning strategy σ' of Eve in the new game G'. Because of the lack of additional choices of Eve in $\mathtt{B}'$ and the fact that the choices of Eve over $\mathtt{B}'$ cannot depend on the choices made by Adam on the sub-boards $\mathcal{F}(\psi)$, we know that σ' corresponds to a pure strategy σ over the board $\mathtt{B}$. We will prove that σ is winning in G. Consider any pure strategy π of Adam over $\mathtt{B}$. Our aim is to prove that $\sigma \cap \pi \in \mathrm{L}(\mathcal{A})$. Assume to the contrary, that $\sigma \cap \pi \notin \mathrm{L}(\mathcal{A})$ and let $\bar{\pi}$ be a positional winning strategy of Adam in the game $\mathrm{G}(\mathcal{A}, \sigma \cap \pi)$. Similarly as above, the pair of strategies π and $\bar{\pi}$ can be combined into one pure strategy π' of Adam over $\mathtt{B}'$.

We will prove that $\sigma' \cap \pi' \notin \mathrm{L}(\mathcal{A}')$, contradicting the assumption that σ' was winning in the game G'. If it was the case that $\sigma' \cap \pi' \in \mathrm{L}(\mathcal{A}')$ then a winning strategy of Eve in $\mathrm{G}(\mathcal{A}', \sigma' \cap \pi')$ would translate into a strategy $\bar{\sigma}$ in $\mathrm{G}(\mathcal{A}, \sigma \cap \pi)$ that would win with the positional strategy $\bar{\pi}$ of Adam. Thus, $\sigma' \cap \pi' \notin \mathrm{L}(\mathcal{A}')$.

G Undecidability of the stochastic value problems

In this section we prove that Problem 14 is undecidable and provide a proof of Theorem 15.

G.1 The word problem for VSNA is undecidable

To prove that Problem 14 is undecidable we will rely on an undecidability result regarding *simple probabilistic automata*, see [13] for details.

▶ **Definition 20.** A probabilistic automaton is called *simple* if every transition probability is in the set $\{0, \frac{1}{2}, 1\}$.

▶ **Theorem 21** ([13])**.** *The problem whether a given simple probabilistic automaton accepts a finite word with probability strictly greater than $\frac{1}{2}$ is undecidable.*

Using the above result, we can prove that Problem 14 is undecidable by the following reduction. First we reduce the alphabet of a given probabilistic automaton $\mathcal{A}'$ to the two-letter alphabet $\{a, b\}$ by a standard technique of encoding letters in binary. We then transform the automaton into a very simple non-deterministic automaton $\mathcal{A}'$ in such a way that

$$\begin{aligned} \forall w \in \{a,b\}^*. \qquad & \mathcal{A} \text{ accepts } w \text{ with probability strictly greater than } \frac{1}{2} \\ & \text{if and only if} \\ & \text{more than half of the runs of } \mathcal{A}' \text{ over } w \text{ is accepting.} \end{aligned} \tag{10}$$

Assume that the set of states of $\mathcal{A}$ is Q. Let the set of states of $\mathcal{A}'$ be $Q \times \{0, 1\}$. The initial state of $\mathcal{A}'$ is $(q_{\mathrm{I}}, 0)$, where q_{I} is the initial state of $\mathcal{A}$.

If the probability of a transition (q, a, q') in $\mathcal{A}'$ is $\frac{1}{2}$ then we just add transitions $\big((q, i), a, (q', i)\big)$ for $i = 0, 1$ to $\mathcal{A}'$. If the probability of a transition (q, a, q') in $\mathcal{A}$ is 1

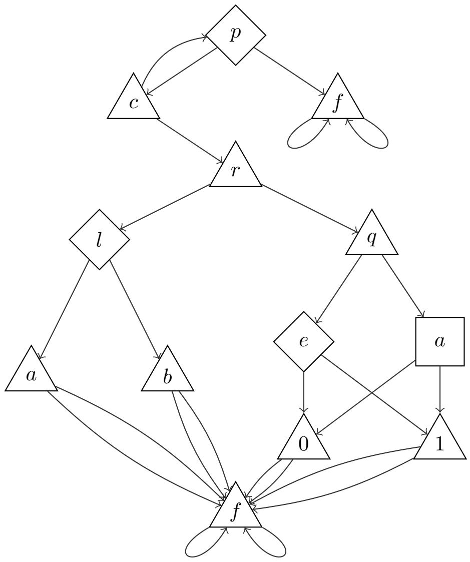

Figure 7 A complete branching board used in the proof of Theorem 15.

then we add transitions $\big((q,i),a,(q',j)\big)$ for $i,j=0,1$ to $\mathcal{A}'$. If the probability of a transition (q,a,q') in $\mathcal{A}$ is 0 then we do not copy this transition to $\mathcal{A}'$.

The accepting states of $\mathcal{A}'$ are of the form (q,i) for q an accepting state of $\mathcal{A}$ and $i \in \{0,1\}$.

Condition (10) is clearly satisfied and $\mathcal{A}'$ is a VSNA. Thus the reduction is completed.

G.2 Proof of Theorem 15

▶ **Theorem 15.** *For every $V \in \{val^{BA}, val^{MA}, val^{ME}, val^{BE}\}$ the value V problem of a regular $\{E, A, \mathcal{B}\}$-branching game is undecidable.*

By duality we can assume that $V \in \{val^{ME}, val^{BE}\}$. Let $\mathtt{B}'$ be the board depicted in Figure 4a with the *Nature*'s position and it's children replaced by the gadget from Figure 4b. The result of this operation can be seen in Figure 7. The idea behind this transformation is the following. If Adam and Eve choose the same vertex in a single instance of the gadget we interpret it as if the *Nature* would have chosen the 0-labelled vertex on the previous board, otherwise we interpret it as choosing the 1-labelled vertex. Since both the players can enforce the $\frac{1}{2}$ probabilities of *Nature*'s choices, the new game has the same value as the previous one, see Claim G.1.

Similarly as in the proof of Theorem 13, a play over the board $\mathtt{B}'$ produces three sequences of the same length:

- the labels $l_0, l_1, \ldots$ with $l_i \in \{a,b\}$ the label of the child of the i-th copy of the vertex labelled l (chosen by Eve),
- the labels $e_0, e_1, \ldots$ with $e_i \in \{0,1\}$ the label of the child of the i-th copy of the vertex labelled e (chosen by Eve),
- the labels $a_0, a_1, \ldots$ with $a_i \in \{0,1\}$ the label of the child of the i-th copy of the vertex labelled a (chosen by Adam).

Let $\vec{e} \otimes \vec{e}$ be the vector obtained by the coordinatewise XOR (i.e. addition modulo 2) of the vectors $\vec{e}$ and $\vec{a}$. Consider the winning condition defined as:

$$L \stackrel{\text{def}}{=} \{t \in \text{plays}(\mathtt{B}) \mid \text{the sequences } \vec{l},\ \vec{e}, \text{ and } \vec{a} \text{ are finite and } \text{run}(\vec{l}, \vec{e} \otimes \vec{a}) \text{ is accepting}\}. \quad (11)$$

Similarly as in the proof of Theorem 15 the above language is regular.

Let $\mathrm{G}' = \langle \mathtt{B}', L' \rangle$. We will conclude the proof of Theorem 13 by showing, that if we take a VSNA $\mathcal{A}$ and construct the two games: G from the proof of Theorem 15 and G′ defined above, then the following claim holds.

▶ **Claim G.1.** *We have*

$$val^{E}_{\mathrm{G}} = val^{ME}_{\mathrm{G}'} = val^{BE}_{\mathrm{G}'}.$$

Clearly $val^{ME}_{\mathrm{G}'} \geq val^{BE}_{\mathrm{G}'}$ by (3).

We start by proving $val^{BE}_{\mathrm{G}'} \geq val^{E}_{\mathrm{G}}$. Consider a pure strategy σ of Eve in G. Clearly this strategy can be extended to a behavioural strategy of Eve in G′ that chooses the successors of the vertices labelled e with equal probability. Now, for every strategy of Adam (no matter if pure, behavioural, or mixed) the probability that a given bit of $\vec{e} \otimes \vec{a}$ is 1 equals $\frac{1}{2}$. Therefore, the probability that $\mathrm{run}(\vec{l}, \vec{e} \otimes \vec{a})$ is accepting is the same in G′ as the probability that $\mathrm{run}(\vec{l}, \vec{n})$ is accepting in G. It implies that $val^{BE}_{\mathrm{G}'} \geq val^{E}_{\mathrm{G}}$.

We finish the proof of Claim G.1 by proving that $val^{E}_{\mathrm{G}} \geq val^{ME}_{\mathrm{G}'}$. For that, we will take a mixed strategy σ'_m of Eve in G′ and prove that that $val^{E}_{\mathrm{G}} \geq val_{\mathrm{G}'}(\sigma'_m)$.

Let us fix a behavioural strategy π_b of Adam that chooses the successors of the vertices labelled a with equal probability. By Lemma 1 and the definition of $val_{\mathrm{G}'}(\sigma'_m)$ we know that

$$val_{\mathrm{G}'}(\sigma'_m, \pi_b) \geq \inf_{\pi \in \Sigma^{BA}_{\mathtt{B}'}} val_{\mathrm{G}'}(\sigma'_m, \pi) = \inf_{\pi \in \Sigma^{A}_{\mathtt{B}'}} val_{\mathrm{G}'}(\sigma'_m, \pi) = val_{\mathrm{G}'}(\sigma'_m).$$

By applying Lemma 1 in Adam's version to the strategy π_b, we get

$$\sup_{\sigma \in \Sigma^{E}_{\mathtt{B}'}} val_{\mathrm{G}'}(\sigma, \pi_b) = \sup_{\sigma \in \Sigma^{ME}_{\mathtt{B}'}} val_{\mathrm{G}'}(\sigma, \pi_b) \geq val_{\mathrm{G}'}(\sigma'_m, \pi_b).$$

It means that for every $\epsilon > 0$ there exists a pure strategy σ' of Eve in G′ such that

$$val_{\mathrm{G}'}(\sigma', \pi_b) \geq val_{\mathrm{G}'}(\sigma'_m) - \epsilon.$$

The strategy σ' induces a pure strategy σ in the game G by forgetting about the choices made in the vertices labelled e. Notice that by the choice of the strategy π_b we know that the probability that for a play resulting in σ' and π_b a given bit of $\vec{e} \otimes \vec{a}$ is 1 equals $\frac{1}{2}$. Therefore, $val_{\mathrm{G}}(\sigma) = val_{\mathrm{G}'}(\sigma', \pi_b)$. It implies that

$$val^{E}_{\mathrm{G}} \geq val_{\mathrm{G}}(\sigma) \geq val_{\mathrm{G}'}(\sigma'_m) - \epsilon,$$

and by the freedom in choosing ϵ we have $val^{E}_{\mathrm{G}} \geq val_{\mathrm{G}'}(\sigma'_m)$.