

PROBABILISTIC MODAL μ -CALCULUS WITH INDEPENDENT PRODUCT*

MATTEO MIO

LIX, Ecole Polytechnique
e-mail address: mio@lix.polytechnique.fr

ABSTRACT. The *probabilistic* modal μ -calculus is a fixed-point logic designed for expressing properties of probabilistic labeled transition systems (PLTS's). Two equivalent semantics have been studied for this logic, both assigning to each state a value in the interval $[0, 1]$ representing the probability that the property expressed by the formula holds at the state. One semantics is *denotational* and the other is a *game semantics*, specified in terms of two-player stochastic parity games.

A shortcoming of the probabilistic modal μ -calculus is the lack of expressiveness required to encode other important temporal logics for PLTS's such as Probabilistic Computation Tree Logic (PCTL). To address this limitation we extend the logic with a new pair of operators: independent product and coproduct. The resulting logic, called *probabilistic modal μ -calculus with independent product*, can encode many properties of interest and subsumes the qualitative fragment of PCTL.

The main contribution of this paper is the definition of an appropriate game semantics for this extended probabilistic μ -calculus. This relies on the definition of a new class of games which generalize standard two-player stochastic (parity) games by allowing a play to be split into concurrent subplays, each continuing their evolution independently. Our main technical result is the equivalence of the two semantics. The proof is carried out in ZFC set theory extended with Martin's Axiom at an uncountable cardinal.

1. INTRODUCTION

The modal μ -calculus ($L\mu$) [18, 30, 5] is a very expressive logic, for expressing properties of labeled transition systems (LTS's), obtained by extending classical propositional modal logic with least and greatest fixed point operators. In the last decade, a lot of research has focused on the study of reactive systems that exhibit some kind of probabilistic behavior, and logics for expressing their properties. Probabilistic labeled transition systems (PLTS's)

1998 ACM Subject Classification: D.2.4, F.3.0, F.4.1.

Key words and phrases: Probabilistic Temporal Logic, Game Semantics, Two-player Stochastic Games, Modal μ -calculus.

* A conference version of this paper, not including proofs, appeared as [24].

This research was partially supported by PhD studentships from LFCS and the IGS at the School of Informatics, University of Edinburgh, and by EPSRC research grant EP-F042043-1. It was completed during the tenure of an ERCIM “Alain Bensoussan” Fellowship, supported by the Marie Curie Co-funding of Regional, National and International Programmes (COFUND) of the European Commission.

[29] are a natural generalization of standard LTS's to the probabilistic scenario, as they allow both non-deterministic and (countable) probabilistic choices. A state s in a PLTS can evolve by non-deterministically choosing one of the *accessible* probability distributions d (over states) and then continuing its execution from the state s' with probability $d(s')$. This combination of non-deterministic choices immediately followed by probabilistic ones, allows the modeling of concurrency, non-determinism and probabilistic behaviors in a natural way. PLTS's can be visualized using graphs labeled with probabilities in a natural way [13, 19, 1]. For example the PLTS depicted in Figure 1 models a system with two states p and q . At the state q no action can be performed. At the state p the system can evolve non-deterministically either to the state q with probability 1 (when the transition $p \xrightarrow{a} d_2$ is chosen) or to the state p with probability $\frac{1}{3}$ and to the state q and with probability $\frac{2}{3}$ (when the transition $p \xrightarrow{a} d_1$ is chosen).




Figure 1: Example of a PLTS

The probabilistic modal μ -calculus ($\text{pL}\mu$), introduced in [27, 14, 10], is a generalization of $\text{L}\mu$ designed for expressing properties of PLTS's. This logic was originally named the *quantitative* μ -calculus, but since other μ -calculus-like logics, designed for expressing properties of non-probabilistic systems, have been given the same name (see, e.g., [11]), we adopt the *probabilistic* adjective. The syntax of the logic $\text{pL}\mu$ coincides with that of the standard μ -calculus. The denotational semantics of $\text{pL}\mu$ [27, 10] generalizes that of $\text{L}\mu$, by interpreting every formula F as a map $\llbracket F \rrbracket : P \rightarrow [0, 1]$, which assigns to each state p a *degree of truth*. In [23], the authors introduce an alternative semantics for the logic $\text{pL}\mu$. This semantics, given in term of two-player stochastic parity games, is a natural generalization of the two-player (non stochastic) game semantics for the logic $\text{L}\mu$ [30]. As in $\text{L}\mu$ games, the two players play a game starting from a configuration $\langle p, F \rangle$, where the objective for Player 1 is to produce a path of configurations along which the outermost fixed point variable X unfolded infinitely often is bound by a greatest fixed point in F . On a configuration of the form $\langle p, G_1 \vee G_2 \rangle$, Player 1 chooses one of the disjuncts G_i , $i \in \{1, 2\}$, by moving to the next configuration $\langle p, G_i \rangle$. On a configuration $\langle p, G_1 \wedge G_2 \rangle$, Player 2 chooses a conjunct G_i and moves to $\langle p, G_i \rangle$. On a configuration $\langle p, \mu X.G \rangle$ or $\langle p, \nu X.G \rangle$ the game evolves to the configuration $\langle p, G \rangle$, after which, from any subsequent configuration $\langle q, X \rangle$ the game again evolves to $\langle q, G \rangle$. On configurations $\langle p, \langle a \rangle G \rangle$ and $\langle p, [a] G \rangle$, Player 1 and 2 respectively choose a transition $p \xrightarrow{a} d$ in the PLTS and move the game to $\langle d, G \rangle$. Here d is a probability distribution (this is the key difference between $\text{pL}\mu$ and $\text{L}\mu$ games) and the configuration $\langle d, G \rangle$ belongs to Nature, the probabilistic agent of the game, who moves on to the next configuration $\langle q, G \rangle$ with probability $d(q)$. This game semantics allows one to interpret formulas as expressing, for each state p , the (limit) probability of a *property*, specified by the formula, holding at the state p . In [23], the equivalence of the denotational and game

semantics for $\text{pL}\mu$ on all finite models, was proven. The result was recently extended to arbitrary models by the present author [26].

Having a complementary game semantics for the logic $\text{pL}\mu$ is of great conceptual importance. In the *quantitative* approach to probabilistic temporal logics, the truth value associated with a formula at a given state is supposed to represent the probability that the property expressed by the formula holds at the state. Since the connectives of $\text{pL}\mu$ can be given other meaningful denotational interpretations different from those considered in [27, 10] (see, e.g., [14]), one naturally seeks an alternative description for the properties associated with formulas going beyond the mere denotational interpretation. The game semantics for $\text{pL}\mu$ provides such an *operational* interpretation in terms of the interactions between the controller (Player 1) and a hostile environment (Player 2) in the context of the stochastic choices occurring in the PLTS (Nature).

However, a shortcoming of the probabilistic μ -calculus is the lack of expressiveness required to encode other important temporal logics for PLTS's, such as Probabilistic Computation Tree Logic (PCTL) of [2]. To address this limitation, we consider an extension of the logic $\text{pL}\mu$ obtained by adding to the syntax of the logic a second conjunction operator (\cdot) called *product* and its De Morgan dual operator called *coproduct* (\odot). We call this extension the *probabilistic modal μ -calculus with independent product*, or just $\text{pL}\mu^\odot$. The denotational semantics of the product operator is defined as $\llbracket F \cdot G \rrbracket(p) = \llbracket F \rrbracket(p) \cdot \llbracket G \rrbracket(p)$, where the product symbol in the right hand side is multiplication on reals. The denotational semantics of the coproduct is defined, by De Morgan duality, as $\llbracket F \odot G \rrbracket(p) = 1 - ((1 - \llbracket F \rrbracket(p)) \cdot (1 - \llbracket G \rrbracket(p)))$. These operators were already considered in [14] as a possible generalization of standard boolean conjunction and disjunction to the lattice $[0, 1]$. Our logic $\text{pL}\mu^\odot$ is novel in containing both ordinary conjunctions and disjunctions (\wedge and \vee) and independent products and coproducts (\cdot and \odot). While giving a denotational semantics to $\text{pL}\mu^\odot$ is straightforward, the major task we undertake in this paper is to extend the game semantics of [23] to the new connectives. The conceptual importance of this kind of result has already been outlined above (see also [25, §1] for an extensive exposition). The game semantics implements the intuition that $H_1 \cdot H_2$ expresses the probability that H_1 and H_2 both hold if verified *independently* of each other.

To capture formally this intuition we introduce a game semantics for the logic $\text{pL}\mu^\odot$ in which independent execution of many instances of the game is allowed. Our games build on those for $\text{pL}\mu$ outlined above. Novelty arises in the game interpretation of the game states $\langle p, H_1 \cdot H_2 \rangle$ and $\langle p, H_1 \odot H_2 \rangle$. When during the execution of the game one of this kind of nodes is reached, the game is split into two concurrent and independent subgames continuing their executions from the states $\langle p, H_1 \rangle$ and $\langle p, H_2 \rangle$ respectively. The difference between the game-interpretation of product and coproduct operators is that on a product configuration $\langle p, H_1 \cdot H_2 \rangle$, Player 1 has to win in both generated sub-games, while on a coproduct configuration $\langle p, H_1 \odot H_2 \rangle$, Player 1 needs to win just one of the two generated sub-games.

To illustrate the main ideas, let us consider the PLTS of figure 2(a) and the $\text{pL}\mu$ formula $F = \langle a \rangle \langle a \rangle \#$ which asserts the possibility of performing two consecutive a -steps. The probability of F being satisfied at p is $\frac{1}{2}$, since after the first a -step, the state $\mathbf{0}$ is reached with probability $\frac{1}{2}$ and no further a -step is possible. Let us consider the $\text{pL}\mu^\odot$ formula $H = \mu X. (F \odot X)$. Figure 2(b) depicts a play in the game starting from the configuration $\langle p, H \rangle$ (fixed-point unfolding steps are omitted). The branching points represent places

whose objective is specified by means of *inner* games, we introduce the notion of *two-player stochastic meta-parity game*.

Our main technical result is the equivalence of the denotational semantics and the game semantics for the logic $\text{pL}\mu^\odot$. As in [26] the proof of equivalence of the two semantics is based on the *unfolding method* of [11]. However there are significant complications, notably, the transfinite inductive characterization of the set of winning branching plays in a given $\text{pL}\mu^\odot$ game (section 6) and the lack of denotational continuity on the free variables taken care by the game-theoretic notion of *robust* Markov branching play (Section 7). Moreover, because of the complexity of the objectives described by means of *inner games*, the proof is carried out in ZFC set theory extended with MA_{\aleph_1} (Martin's Axiom at \aleph_1) and therefore our result is at least consistent with ZFC. We leave open the question of whether our result is provable in ZFC alone; we do not know if this is possible even restricting the equivalence problem to finite models.

The rest of the paper is organized as follows. In Section 2 we discuss the required mathematical background. In Section 3 we define the syntax and the denotational semantics of the logic $\text{pL}\mu^\odot$. In Section 4 the class of stochastic tree games, and its sub-class given by two-player stochastic meta-parity games, are introduced in detail. In Section 5 the game semantics of the logic $\text{pL}\mu^\odot$ is defined in terms of two-player stochastic meta-parity games. In Section 6 we provide a transfinite inductive characterization of the winning set of the game associated with a formula $\mu X.F$. In Section 7 we introduce the technical notion of robust Markov branching play. In Section 8 we prove the equivalence of the two semantics, our main result. Conclusions and directions for future research are presented in Section 9.

2. MATHEMATICAL BACKGROUND

Definition 2.1. A (discrete) probability distribution d over a set X is a function $d: X \rightarrow [0, 1]$ such that $\sum_{x \in X} d(x) = 1$. The *support* of d , denoted by $\text{supp}(d)$, is defined as the (necessarily countable) set $\{x \in X \mid d(x) > 0\}$. We denote with $\mathcal{D}(X)$ the set of probability distributions over X . We denote with δ_x , for $x \in X$, the unique probability distribution such that $\text{supp}(\delta_x) = \{x\}$.

Definition 2.2 (PLTS [29]). Given a countable set L of labels, a *Probabilistic Labeled Transition System* is a pair $\mathcal{L} = \langle P, \{\xrightarrow{a}\}_{a \in L} \rangle$, where P is a countable set of states and $\xrightarrow{a} \subseteq P \times \mathcal{D}(P)$, for every $a \in L$. In this paper we restrict our attention to those PLTS's such that for every $p \in P$ and every $a \in L$, the set $\{d \mid p \xrightarrow{a} d\}$ is countable. We refer to the countable set $\bigcup_{a \in L} \bigcup_{p \in P} \{d \mid p \xrightarrow{a} d\}$, denoted by $\mathcal{D}(\mathcal{L})$, as the set of probability distributions of the PLTS. We say that a PLTS \mathcal{L} is *not-probabilistic*, or just a LTS, if every probability distribution $d \in \mathcal{D}(\mathcal{L})$ is of the form δ_p , for some $p \in P$.

Given a set X , we denote with 2^X the set of all subsets $Y \subseteq X$. Given a complete lattice (L, \leq) , we denote with $\bigsqcup: 2^L \rightarrow L$ and $\bigsqcap: 2^L \rightarrow L$ the operations of join and meet respectively. A function $f: L \rightarrow L$ is *monotone* if $x \leq y$ implies $f(x) \leq f(y)$, for every $x, y \in L$. The set of fixed points of any monotone function $f: L \rightarrow L$, ordered by \leq , is a non-empty complete lattice [32]. We denote with $\text{lfp}(f)$ and $\text{gfp}(f)$ the least and the greatest fixed points of f , respectively.

Theorem 2.3 (Knaster–Tarski [32]). *Let (L, \leq) be a complete lattice and $f: L \rightarrow L$ a monotone function. The following equalities hold:*

- (1) $\text{lfp}(f) = \bigsqcup_{\alpha} f^{\alpha}$, where $f^{\alpha} = \bigsqcup_{\beta < \alpha} f(f^{\beta})$,
 (2) $\text{gfp}(f) = \bigsqcap_{\alpha} f_{\alpha}$, where $f_{\alpha} = \bigsqcap_{\beta < \alpha} f(f_{\beta})$,
 where the greek letters α and β range over ordinals.

The closed real interval $[0, 1]$, with its standard order \leq , is a (distributive) complete lattice [25] with $\bigsqcup X = \sup X$ and $\bigsqcap X = \inf X$, for every $X \subseteq [0, 1]$. When $X = \{x, y\}$ we simply have $x \sqcup y = \max\{x, y\}$ and $x \sqcap y = \min\{x, y\}$. The involutive map $x \mapsto 1 - x$ is clearly order-reversing, thus the structure $([0, 1], \leq, \lambda x. 1 - x)$ constitutes a complete De Morgan algebra. We shall consider the operation of product (\cdot) , i.e., standard multiplication on $[0, 1]$, and its De Morgan dual, called coproduct (\odot) , defined as: $x \odot y \stackrel{\text{def}}{=} 1 - ((1 - x) \cdot (1 - y))$. Equivalently, $x \odot y = x + y - xy$. Both operations are commutative, associative, monotone and extend uniquely to infinitary operations of type $\prod, \coprod : [0, 1]^{\mathbb{N}} \rightarrow [0, 1]$ as follows:

$$\prod \{x_n\}_{n \in \mathbb{N}} = \prod_{m \in \mathbb{N}} x_0 \cdot \dots \cdot x_m \quad \text{and} \quad \coprod \{x_n\}_{n \in \mathbb{N}} = \prod_{m \in \mathbb{N}} x_0 \odot \dots \odot x_m.$$

We shall make use of the following fast growing function on the natural numbers to approximate infinite (co)products.

Definition 2.4. The function $\# : \mathbb{N} \rightarrow \mathbb{N}$ is defined as $\#(n) \stackrel{\text{def}}{=} 2^{2^n + 1}$.

Proposition 2.5. The following inequalities hold for every $\varepsilon \in (0, 1]$ and $\vec{x}_n \in [0, 1]^{\mathbb{N}}$:

$$\prod_{n \in \mathbb{N}} \left(x_n + \frac{\varepsilon}{\#(n)}\right) \leq \left(\prod_{n \in \mathbb{N}} x_n\right) + \varepsilon \quad \text{and} \quad \prod_{n \in \mathbb{N}} \left(x_n - \frac{\varepsilon}{\#(n)}\right) \geq \left(\prod_{n \in \mathbb{N}} x_n\right) - \varepsilon$$

and, similarly, if infinitary products are replaced by infinitary coproducts.

Proof. Both inequalities are easily proved by routine methods. We refer to Lemma 2.2.10 of [25] for a detailed proof. \square

In the following we assume standard notions of basic topology and basic measure theory. We refer to [17] and [31] as standard references to these topics. The author's PhD thesis [25, Chapter 2] provides a succinct introduction to the necessary material. The topological spaces we consider will always be 0-dimensional Polish spaces.

Adopting standard notation, we denote with $\mathbf{\Gamma}_{\alpha}^0$, for $\Gamma \in \{\Sigma, \Pi, \Delta\}$ and α a countable ordinal greater than 0, the corresponding class of subsets (of a Polish space) in the Borel hierarchy. In particular, $\mathbf{\Delta}_1^0$, $\mathbf{\Sigma}_1^0$ and $\mathbf{\Pi}_1^0$ denote the collections of clopen, open and closed sets respectively.

Definition 2.6. For each $n \geq 1$ the *projective classes* $\mathbf{\Sigma}_n^1(X)$, $\mathbf{\Pi}_n^1(X)$, $\mathbf{\Delta}_n^1(X)$ of subsets of a Polish space X are defined as follows:

$$\begin{aligned} \mathbf{\Sigma}_{n+1}^1 &= \left\{ A \subseteq X \mid A = \{x \mid \exists y. (x, y) \in B\} \text{ for some } B \in \mathbf{\Pi}_n^1(X \times Y) \right\} \\ \mathbf{\Pi}_{n+1}^1 &= \left\{ A \subseteq X \mid (X \setminus A) \in \mathbf{\Sigma}_{n+1}^1(X) \right\} \\ \mathbf{\Delta}_n^1 &= \mathbf{\Sigma}_n^1 \cap \mathbf{\Pi}_n^1 \end{aligned}$$

where Y is a Polish space, and $X \times Y$ is equipped with the product topology. It is a well-known result of Susin (see, e.g., Theorem 14.11 in [17]) that $\mathbf{\Delta}_1^1$ is the collection of Borel sets. Thus the class of projective sets is defined, starting from the Borel sets, by iterating the operations of projection and complementation.

Definition 2.7. Given a Polish space X and a subset $A \subseteq X$, we say that A is *universally measurable* if it is μ -measurable for every Borel probability measure on X . It follows that

the collection of universally measurable subsets of X , denoted by $\text{UM}(X)$, forms a σ -algebra. Given a measurable space (Y, \mathcal{S}) and a function $f : X \rightarrow Y$, we say that f is *universally measurable* if $f^{-1}(S) \in \text{UM}(X)$ for all $S \in \mathcal{S}$.

The following facts will be useful.

Theorem 2.8. *Let X be a Polish space. The following assertions hold:*

- (1) $\Sigma_1^1(X) \subseteq \text{UM}(X)$, and if X is uncountable the inclusion is strict.
- (2) If $A \in \text{UM}(Y)$ and $f : X \rightarrow Y$ is universally measurable, then $f^{-1}(A) \in \text{UM}(X)$.

It is not possible, in ZFC alone, to show that universally measurable sets extend any further up in the projective hierarchy.

Theorem 2.9. *Let X be a Polish space. If ZFC is consistent the following assertions hold:*

- (1) $\text{ZFC} \not\vdash \Delta_2^1(X) \subseteq \text{UM}(X)$,
- (2) $\text{ZFC} \not\vdash \Delta_2^1(X) \subsetneq \text{UM}(X)$.

Proof. The result of the first assertion is due to Kurt Gödel, see, e.g., Corollary 25.28 of [15]. A proof of the second assertion can be found in, e.g., [22]. \square

Thus, it is not decidable in ZFC if every set in $\Delta_2^1(X)$ is universally measurable. As we shall see in later sections, the winning sets of $\text{pL}\mu^\odot$ games are, in general, $\Delta_2^1(X)$ sets. To deal with the associated measure theoretic issues, we will carry our main result in $\text{ZFC} + \text{MA}_{\aleph_1}$ set theory, where MA_{\aleph_1} is the so-called *Martin's Axiom at \aleph_1* , the first uncountable cardinal. We often identify \aleph_1 with the least uncountable ordinal ω_1 . We refer to [22, 25] for a description of the axiom MA_{\aleph_1} and some of its set-theoretic consequences. Here we list those that are relevant to our work.

Theorem 2.10 ([22]). *Let X be a Polish space, μ be a Borel probability measure on X and $\{A_\alpha\}_{\alpha < \omega_1}$ a collection of μ -measurable subsets of X , where ω_1 is the least uncountable ordinal. The following assertions hold in $\text{ZFC} + \text{MA}_{\aleph_1}$:*

- (1) $\Delta_2^1(X) \subseteq \text{UM}(X)$,
- (2) ω_1 -completeness: $\bigcup_{\alpha < \omega_1} A_\alpha$ is a μ -measurable set,
- (3) ω_1 -continuity: $\mu(\bigcup_{\alpha < \omega_1} A_\alpha) = \bigsqcup_{\alpha < \omega_1} \mu(A_\alpha)$.

Note that, since singleton sets are always closed (hence measurable) in Polish spaces, it follows immediately from Theorem 2.10 that Martin's Axiom at \aleph_1 implies $2^{\aleph_0} \neq \aleph_1$, i.e., the negation of *Continuum Hypothesis* [15]. As a matter of fact, Martin's Axiom was introduced by the authors as a possible set-theoretic alternative to the Continuum Hypothesis [22].

3. THE LOGIC $\text{pL}\mu^\odot$

Given a set Var of propositional variables ranged over by the letters X, Y and Z , and a set of labels L ranged over by the letters a, b and c , the formulas of the logic $\text{pL}\mu^\odot$ are defined by the following grammar:

$$F, G ::= X \mid F \wedge G \mid F \vee G \mid F \cdot G \mid F \odot G \mid [a]F \mid \langle a \rangle F \mid \nu X.F \mid \mu X.F$$

which extends the syntax of the probabilistic modal μ -calculus ($\text{pL}\mu$) with a new kind of conjunction (\cdot) and disjunction (\odot) operators called *product* and *coproduct* respectively. As usual the operators $\nu X.F$ and $\mu X.F$ bind the variable X in F . A formula is *closed* if it has no *free* variables.

Definition 3.1 (Subformula). We define the set $Sub(F)$ by case analysis on F as follows: $Sub(X) \stackrel{\text{def}}{=} \{X\}$, $Sub([a] F) = \{[a] F\} \cup Sub(F)$, $Sub(F_1 \wedge F_2) = \{F_1 \wedge F_2\} \cup Sub(F_1) \cup Sub(F_2)$, $Sub(F_1 \cdot F_2) = \{F_1 \cdot F_2\} \cup Sub(F_1) \cup Sub(F_2)$ and $Sub(\nu X.F) = \{\nu X.F\} \cup Sub(F)$. The cases for the connectives $\langle a \rangle$, \vee , \odot and μX are defined as for their duals. We say that G is a *subformula* of F if $G \in Sub(F)$.

Given a PLTS $\langle P, \{\xrightarrow{a}\}_{a \in L} \rangle$ we denote with $[0, 1]^P$ the complete lattice of functions from P to the real interval $[0, 1]$ with the pointwise order. A function $\rho : Var \rightarrow [0, 1]^P$ is called a $[0, 1]$ -valued *interpretation*, or just an interpretation, of the variables. Given a function $f : P \rightarrow [0, 1]$ we denote with $\rho[f/X]$ the interpretation that assigns f to the variable X , and $\rho(Y)$ to all other variables Y .

The denotational semantics $\llbracket F \rrbracket_\rho : P \rightarrow [0, 1]$ of the $\text{pL}\mu^\odot$ formula F , under the interpretation ρ , is defined by structural induction on F as follows:

$$\begin{aligned}
\llbracket X \rrbracket_\rho(p) &= \rho(X)(p) \\
\llbracket G \vee H \rrbracket_\rho(p) &= \llbracket G \rrbracket_\rho(p) \sqcup \llbracket H \rrbracket_\rho(p) \\
\llbracket G \wedge H \rrbracket_\rho(p) &= \llbracket G \rrbracket_\rho(p) \sqcap \llbracket H \rrbracket_\rho(p) \\
\llbracket G \odot H \rrbracket_\rho(p) &= \llbracket G \rrbracket_\rho(p) \odot \llbracket H \rrbracket_\rho(p) \\
\llbracket G \cdot H \rrbracket_\rho(p) &= \llbracket G \rrbracket_\rho(p) \cdot \llbracket H \rrbracket_\rho(p) \\
\llbracket \langle a \rangle G \rrbracket_\rho(p) &= \bigsqcup_{p \xrightarrow{a} d} \left(\sum_{q \in \text{supp}(d)} d(q) \cdot \llbracket G \rrbracket_\rho(q) \right) \\
\llbracket [a] G \rrbracket_\rho(p) &= \prod_{p \xrightarrow{a} d} \left(\sum_{q \in \text{supp}(d)} d(q) \cdot \llbracket G \rrbracket_\rho(q) \right) \\
\llbracket \mu X.G \rrbracket_\rho(p) &= \text{lfp} \left(\lambda f. (\llbracket G \rrbracket_{\rho[f/X]}) \right)(p) \\
\llbracket \nu X.G \rrbracket_\rho(p) &= \text{gfp} \left(\lambda f. (\llbracket G \rrbracket_{\rho[f/X]}) \right)(p)
\end{aligned}$$

where the symbols \cdot and \odot on the right hand side denote the operations of product and coproduct on $[0, 1]$, respectively. It is easy to verify that the interpretation assigned to every $\text{pL}\mu^\odot$ operator is monotone, thus the existence of the least and greatest fixed points is guaranteed by the Knaster–Tarski theorem.

The interpretation of the connectives of the logic $\text{pL}\mu^\odot$ (and its fragment $\text{pL}\mu$) resembles the corresponding ones for $\text{L}\mu$. Both operations $\{\sqcup, \odot\}$, when restricted to the two element set $\{0, 1\}$ act as ordinary boolean disjunction. Similarly, the operations $\{\sqcap, \cdot\}$ restricts to ordinary boolean conjunction. For what concerns the semantics associated with the modal operators, since in PLTS's transitions lead to probability distributions over states, rather than states, the most natural way to interpret the meaning of a formula G at a probability distribution d is to consider the *expected probability* of the formula G holding at a state q randomly drawn in accordance with d , and this is formalized by the weighted sums in the definition above.

Remark 3.2. As it is common practice when dealing with fixed point logics such as the modal μ -calculus, we presented the syntax of $\text{pL}\mu^\odot$ in *positive form*, i.e., without including a negation operator. This simplifies the presentation of the denotational semantics because all formulas in positive form are interpreted as monotone functions. A negation operator on (closed) $\text{pL}\mu^\odot$ formulas can be defined by induction on the structure of the formula, by exploiting the dualities between the connectives of the logic, in such a way that $\llbracket \neg F \rrbracket_\rho(p) = 1 - \llbracket F \rrbracket_\rho(p)$, for all formulas F and states p . We omit the routine details.

As anticipated in the introduction, the main reason for extending the logic $pL\mu$ with the new pair of connectives of product and coproduct is to get a richer and more expressive logic capable of encoding the *qualitative threshold modalities* defined as follows.

Definition 3.3. Given a $pL\mu^\odot$ formula F , we define the macro formulas $\mathbb{P}_{>0}F$ and $\mathbb{P}_{=1}F$ as follows: $\mathbb{P}_{>0}F \stackrel{\text{def}}{=} \mu X.(F \odot X)$ and $\mathbb{P}_{=1}F \stackrel{\text{def}}{=} \nu X.(F \cdot X)$, where X is not free in F .

The following lemma captures the denotational semantics of the qualitative threshold modalities.

Lemma 3.4. Given a PLTS $\mathcal{L} = \langle P, \{\xrightarrow{a}\}_{a \in L} \rangle$, a $[0, 1]$ -interpretation of the variables $\rho \in \text{Var} \rightarrow [0, 1]^P$ and a $pL\mu^\odot$ formula F , the following assertions hold:

$$\llbracket \mathbb{P}_{>0}F \rrbracket_\rho(p) = \begin{cases} 1 & \text{if } \llbracket F \rrbracket_\rho(p) > 0 \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad \llbracket \mathbb{P}_{=1}F \rrbracket_\rho(p) = \begin{cases} 1 & \text{if } \llbracket F \rrbracket_\rho(p) = 1 \\ 0 & \text{otherwise} \end{cases}$$

for every state $p \in P$.

Proof. The map $x \mapsto \lambda \odot x$, for a fixed $\lambda \in [0, 1]$, has 1 as unique fixed point when $\lambda > 0$, and 0 as the least fixed point when $\lambda = 0$. Similarly for the map $x \mapsto \lambda \cdot x$. The result then follows trivially. \square

It is easy to verify that the derived qualitative threshold modalities, if taken as primitives, are De Morgan duals. As an immediate consequence of Lemma 3.4 we have the following fact.

Proposition 3.5. The denotational interpretation of an open $pL\mu^\odot$ formula F is, in general, not continuous in the free variables, i.e., the denotation of a formula, seen as a function of type $(\text{Var} \rightarrow [0, 1]^P) \rightarrow [0, 1]^P$, where $(\text{Var} \rightarrow [0, 1]^P)$ is ordered pointwise, does not preserve countable \sqsubseteq -increasing chains.

Proof. Consider the formula $\mathbb{P}_{=1}X = \nu Y.(X \cdot Y)$ having just one free variable X . We have that $\llbracket \mathbb{P}_{=1}X \rrbracket_\rho(p) = 0$ if $\rho(X)(p) < 1$ and $\llbracket \mathbb{P}_{=1}X \rrbracket_\rho^\mathcal{L}(p) = 1$ if $\rho(X)(p) = 1$. \square

This contrasts with the fact that the denotations of $pL\mu$ formulas (i.e., formulas without occurrences of (co)products operators) when interpreted over finite PLTS's, are continuous in the free variables (see, e.g., Appendix C of [23] for a proof of this fact, and [25, §3.3.2.3] for a discussion about this phenomenon). It then follows that the qualitative threshold modalities are not expressible in $pL\mu$. Thus, $pL\mu^\odot$ is a strictly more expressive logic than $pL\mu$, as previously claimed.

The following theorem summarizes some expressivity results about the logic $pL\mu^\odot$. Since the focus of this paper is primarily the study of a game semantics for $pL\mu^\odot$, we just refer to Theorem 7.2.16 and Proposition 7.2.5 in [25] for detailed proofs.

Theorem 3.6. The following propositions hold. *s.*

- (1) the logic $pL\mu^\odot$ can encode the qualitative fragment of the logic PCTL of [2],
- (2) there are (closed) $pL\mu^\odot$ formulas that can be satisfied only by infinite PLTS's. Thus $pL\mu^\odot$ does not satisfy the (expected adaptation of) the finite model property,
- (3) there are (closed) $pL\mu^\odot$ formulas satisfiable by some PLTS's but not satisfiable by any non-probabilistic PLTS (see Definition 2.2),

where a PLTS $\mathcal{L} = \langle P, \{\xrightarrow{a}\}_{a \in L} \rangle$ satisfies a (closed) formula F , if there is some $p \in P$ such that $\llbracket F \rrbracket(p) = 1$.

4. STOCHASTIC TREE GAMES

In this (unavoidably long) section we introduce a new class of games which we call *two-player stochastic tree games*, or just $2\frac{1}{2}$ -player *tree games*. Stochastic tree games generalizes standard two-player turn-based stochastic games (see, e.g., [7], [8]) by allowing a new class of *branching states* on which the execution of the game is split in independent concurrent subgames. Formally, stochastic tree games are infinite duration games played by Player 1, Player 2 and a third probabilistic agent named *Nature*, on a *game arena* $\mathcal{A} = \langle (S, E), (S_1, S_2, S_N, B), \pi \rangle$, where (S, E) is a directed graph with countable set of vertices S and transition relation E , (S_1, S_2, S_N, B) is a partition of S and $\pi: S_N \rightarrow \mathcal{D}(S)$. The states in S_1 , S_2 , S_N and S_B are called *Player 1 states*, *Player 2 states*, *probabilistic states* and *branching states* respectively. We denote with $E(s)$, for $s \in S$, the set $\{s' \mid (s, s') \in E\}$. As a technical constraint, we require¹ that $\text{supp}(\pi(s)) \subseteq E(s)$, for every $s \in S_N$.

Definition 4.1 (Paths in \mathcal{A}). We denote with \mathcal{P}^ω and $\mathcal{P}^{<\omega}$ the sets of infinite and finite paths in \mathcal{A} . Given a finite path $\vec{s} \in \mathcal{P}^{<\omega}$ we denote with $\text{first}(\vec{s})$ and $\text{last}(\vec{s})$ the first and last state of \vec{s} , respectively. Given a finite path \vec{s} and a (finite or infinite) path \vec{t} we write $\vec{s} \triangleleft \vec{t}$ if \vec{s} is a (not necessarily proper) prefix of \vec{t} and, provided that $\text{first}(\vec{t}) \in E(\text{last}(\vec{s}))$, we write $\vec{s}.\vec{t}$ for the concatenation of the two paths. We denote with \mathcal{P}^t the set of finite paths ending in a terminal state, i.e., the set of paths \vec{s} such that $E(\text{last}(\vec{s})) = \emptyset$. Similarly, we denote with $\mathcal{P}_1^{<\omega}$ and $\mathcal{P}_2^{<\omega}$ the sets of finite paths ending in a state in S_1 and S_2 , respectively. We denote with \mathcal{P} the set $\mathcal{P}^\omega \cup \mathcal{P}^t$ and we refer to this set as the set of *completed paths* in \mathcal{A} . Given a finite path $\vec{s} \in \mathcal{P}^{<\omega}$, we denote with $O_{\vec{s}}$ the set of all completed paths having \vec{s} as prefix. We consider the standard topology on \mathcal{P} where the countable basis for the open sets is given by the clopen sets $O_{\vec{s}}$, for $\vec{s} \in \mathcal{P}^{<\omega}$. This is a 0-dimensional Polish space [25].

Definition 4.2 (Tree in \mathcal{A}). A *tree* in the arena \mathcal{A} is a collection $T = \{\vec{s}_i\}_{i \in I}$ of finite paths $\vec{s}_i \in \mathcal{P}^{<\omega}$, such that

- (1) T is down-closed: if $\vec{s} \in T$ and $\vec{t} \triangleleft \vec{s}$ then $\vec{t} \in T$.
- (2) T has a root: there exists exactly one finite path $\{s\}$ of length one in T . The state s , denoted by $\text{root}(T)$, is called the root of the tree T .

We consider the nodes \vec{s} of T as labeled by the last function.

Definition 4.3 (Uniquely and fully branching nodes of a tree). A node \vec{s} in a tree T is said to be *uniquely branching* in T if either $E(\text{last}(\vec{s})) = \emptyset$ or \vec{s} has a unique child in T . Similarly, \vec{s} is *fully branching* in T if, for every $s \in E(\text{last}(\vec{s}))$, it holds that $\vec{s}.s \in T$.

An outcome of the game in \mathcal{A} , which we call a *branching play*, is a possibly infinite tree T in \mathcal{A} defined as follows:

Definition 4.4 (Branching play in \mathcal{A}). A *branching play* in the arena \mathcal{A} is a tree T in \mathcal{A} such that, for every node $\vec{s} \in T$ the following conditions holds:

- (1) If $\text{last}(\vec{s}) \in S_1 \cup S_2 \cup S_N$ then \vec{s} branches uniquely in T .
- (2) If $\text{last}(\vec{s}) \in B$ then \vec{s} branches fully in T .

We denote with \mathcal{BP} the set of branching plays T in the arena \mathcal{A} .

¹The constraint $\text{supp}(\pi(s)) = E(s)$, which might look more natural, is unnecessarily restrictive since one just one want to impose that Nature's choices always belong to the set of successor states. Our relaxed constraint will be technically convenient, see, e.g., Footnote 3.

A branching play T represents a possible execution of the game from the state s labeling the root of T . The nodes of T with more than one child are all labeled with a state $s \in B$ and are the branching points of the game. Their children represent the many independent instances of play generated at the branching point.

Definition 4.5 (Topology on \mathcal{BP}). Given a finite tree F in \mathcal{A} , we denote with $O_F \subseteq \mathcal{BP}$ the set of all branching plays T such that $F \subseteq T$. We fix the topology on \mathcal{BP} , where the basis for the open sets is given by the clopen sets O_F , for every branching-play prefix F . It is routine to show that this is a 0-dimensional Polish space [25].

As usual when working with *stochastic* games, it is useful to look at the possible outcomes of a play up-to the behavior of Nature. In the context of standard two-player stochastic games this amounts to considering Markov chains. In our setting the following definition of Markov branching play is natural:

Definition 4.6 (Markov branching play in \mathcal{A}). A *Markov branching play* in \mathcal{A} is a tree M in \mathcal{A} such that for every node $\vec{s} \in M$, the following conditions hold:

- (1) If $\text{last}(\vec{s}) \in S_1 \cup S_2$ then \vec{s} branches uniquely in T .
- (2) If $\text{last}(\vec{s}) \in S_N \cup B$ then \vec{s} branches fully in T .

A Markov branching play is similar to a branching play except that probabilistic choices of Nature have not been resolved.

Definition 4.7 (Probability measure \mathbb{P}_M). Every Markov branching play M determines a probability assignment $\mathbb{P}_M(O_F)$ to every basic clopen set $O_F \subseteq \mathcal{BP}$, for F a finite tree in \mathcal{A} (we can assume that every node $\vec{s} \in F$ such that $\text{last}(\vec{s}) \in S_N$ has a unique child in F , since otherwise $O_F = \emptyset$), defined as follows:

$$\mathbb{P}_M(O_F) \stackrel{\text{def}}{=} \begin{cases} \prod \{\pi(s)(s') \mid \vec{s}.s.s' \in F \wedge s \in S_N\} & \text{if } F \subseteq M \\ 0 & \text{otherwise} \end{cases}$$

Such an assignment extends, by Carathéodory's Extension Theorem [31, 25], to a unique Borel probability measure on \mathcal{BP} , whence to a complete probability measure, also denoted by \mathbb{P}_M .

It is the definition above that implements the *probabilistic independence* of the sub-branching plays rooted at some branching node.

Definition 4.8 (Two-player stochastic tree game). A *two-player stochastic tree game* (or a $2\frac{1}{2}$ -player tree game) is given by a pair $\langle \mathcal{A}, \Phi \rangle$, where \mathcal{A} is a stochastic tree game arena as described above, and $\Phi \subseteq \mathcal{BP}$, which is the *objective* or *winning set* for Player 1, is a universally measurable set of branching plays in \mathcal{A} .

Definition 4.9 (Expected value of a Markov branching play). Let $\langle \mathcal{A}, \Phi \rangle$ be a $2\frac{1}{2}$ -player tree game, and M a Markov branching play in \mathcal{A} . We define the *expected value* of M as follows: $\mathbf{E}(M) = \mathbb{P}_M(\Phi)$. The value $\mathbf{E}(M)$ should be understood as the probability for Player 1 to win the probabilistic play represented by M .

As usual in game theory, players' moves are determined by strategies.

Definition 4.10 (Deterministic strategies). An (*unbounded memory deterministic*) strategy σ_1 for Player 1 in \mathcal{A} is defined as a function $\sigma_1 : \mathcal{P}_1^{<\omega} \rightarrow S \cup \{\bullet\}$ such that $\sigma_1(\vec{s}) \in E(\text{last}(s))$ if $E(\text{last}(\vec{s})) \neq \emptyset$ and $\sigma_1(\vec{s}) = \bullet$ otherwise. Similarly a *strategy* σ_2 for Player 2 is defined as

a function $\sigma_2: \mathcal{P}_2^{<\omega} \rightarrow S \cup \{\bullet\}$. A pair $\langle \sigma_1, \sigma_2 \rangle$ of strategies, one for each player, is called a *strategy profile* and determines the behaviors of both players.

Note that the above definition of strategy captures the intended behavior of the game: both players, when acting on a given instance of the game, know all the history of the actions happened on that subgame, but have no knowledge of the evolution of the other *independent* parallel subgames.

Definition 4.11 ($M_{\sigma_1, \sigma_2}^{s_0}$). Given an initial state $s_0 \in S$ and a strategy profile $\langle \sigma_1, \sigma_2 \rangle$ a unique Markov branching play $M_{\sigma_1, \sigma_2}^{s_0}$ is determined:

- (1) the root of M is labeled with s_0 ,
- (2) for every $\vec{s} \in M_{\sigma_1, \sigma_2}^{s_0}$, if $\text{last}(\vec{s}) = s$ with $s \in S_1$ not a terminal state, then the unique child of \vec{s} in $M_{\sigma_1, \sigma_2}^{s_0}$ is $\vec{s}.(\sigma_1(\vec{s}))$,
- (3) for every $\vec{s} \in M_{\sigma_1, \sigma_2}^{s_0}$, if $\text{last}(\vec{s}) = s$ with $s \in S_2$ not a terminal state, then the unique child of \vec{s} in $M_{\sigma_1, \sigma_2}^{s_0}$ is $\vec{s}.(\sigma_2(\vec{s}))$.

This specifies uniquely $M_{\sigma_1, \sigma_2}^{s_0}$ because Markov branching plays branch fully on probabilistic and branching states.

Definition 4.12 (Upper and lower values of a $2\frac{1}{2}$ -player tree game). Let $\mathcal{G} = \langle \mathcal{A}, \Phi \rangle$ be a $2\frac{1}{2}$ -player tree game. We define the lower and upper values of \mathcal{G} on the state s , denoted by $\text{VAL}_{\downarrow}^s(\mathcal{G})$ and $\text{VAL}_{\uparrow}^s(\mathcal{G})$ respectively, as follows:

$$\text{VAL}_{\downarrow}^s(\mathcal{G}) = \bigsqcup_{\sigma_1} \bigcap_{\sigma_2} \mathbf{E}(M_{\sigma_1, \sigma_2}^s) \quad \text{VAL}_{\uparrow}^s(\mathcal{G}) = \bigcap_{\sigma_2} \bigsqcup_{\sigma_1} \mathbf{E}(M_{\sigma_1, \sigma_2}^s)$$

$\text{VAL}_{\downarrow}^s(\mathcal{G})$ represents the limit probability of Player 1 winning, when the game begins at s , by choosing his strategy σ_1 first and then letting Player 2 pick an appropriate counter strategy σ_2 . Similarly $\text{VAL}_{\uparrow}^s(\mathcal{G})$ represents the limit probability of Player 1 winning, when the game begins at s , by first letting Player 2 choose a strategy σ_2 and then picking an appropriate counter strategy σ_1 . Clearly, for every s , the following inequality holds: $\text{VAL}_{\downarrow}^s(\mathcal{G}) \leq \text{VAL}_{\uparrow}^s(\mathcal{G})$. In the special case (not true in general) that this inequality is an equality, we say that the game \mathcal{G} is *determined* at s .

Definition 4.13 (ε -optimal strategies). Let $\mathcal{G} = \langle \mathcal{A}, \Phi \rangle$ be a $2\frac{1}{2}$ -player tree game. We say that a strategy σ_1 for Player 1 in \mathcal{G} is ε -optimal, for $\varepsilon \geq 0$, if for every state s , the following inequality holds:

$$\bigcap_{\sigma_2} \mathbf{E}(M_{\sigma_1, \sigma_2}^s) \geq \text{VAL}_{\downarrow}^s(\mathcal{G}) - \varepsilon.$$

Similarly we say that a deterministic strategy σ_2 for Player 2 in \mathcal{G} is ε -optimal, for $\varepsilon \geq 0$, if for every state s , the following inequality holds:

$$\bigsqcup_{\sigma_1} \mathbf{E}(M_{\sigma_1, \sigma_2}^s) \leq \text{VAL}_{\uparrow}^s(\mathcal{G}) + \varepsilon.$$

Clearly ε -optimal (mixed and deterministic) strategies for Player 1 and Player 2 always exist for every $\varepsilon > 0$, but not necessarily so for $\varepsilon = 0$.

Remark 4.14. Observe that a $2\frac{1}{2}$ -player tree game without branching states is just an ordinary two-player turn-based stochastic game (see, e.g., [8], [25]): the set of branching plays is homeomorphic to the set of completed paths, the notion of Markov branching play collapses to the standard notion of Markov chain, and strategies are maps from finite paths to successor states. Thus, as previously claimed, two-player stochastic tree games constitute a generalization of ordinary two-player stochastic games.

Notwithstanding our primary interest in an appropriate game semantics for $\text{pL}\mu^\odot$, we highlighting here that the simple form of *partial information* implemented in tree games (players when acting on a given subgame are not aware of what happens in the other independent subgames) is surprisingly powerful. For instance the class of Blackwell games [3, 21] can be encoded as two-players tree games (Theorem 4.2.18 in [25]). Moreover, interesting open problems in the field, such as that of *qualitative determinacy* of stochastic games (see, e.g., [6]), can be formulated as appropriate determinacy problems for 2-player tree games (see, [25, §4.4]). We refer to the author's PhD thesis [25] for an extensive introduction and analysis of tree games.

4.1. Two-player stochastic meta-parity games. In this subsection we identify a class of two-player stochastic tree games, called $2\frac{1}{2}$ -player meta-parity games, which will be used to give an appropriate game semantics to the logic $\text{pL}\mu^\odot$.

A $2\frac{1}{2}$ -player meta-parity game is a $2\frac{1}{2}$ -player tree game $\mathcal{G} = \langle \mathcal{A}, \Phi \rangle$, which we refer to as the *outer game*, in which every branching play T is itself interpreted as a (ordinary) two-player parity game G_T , which we refer to as the *inner game associated with T* , and whose objective Φ is defined as the set of branching plays T for which Player 1 has a winning strategy in G_T .

We start formalizing this notion with the following definitions.

Definition 4.15. A *parity assignment* Ω for a two-player stochastic tree game arena $\mathcal{A} = \langle (S, E), (S_1, S_2, S_N, B), \pi \rangle$ is a function $\Omega : S \rightarrow \mathbb{N}$ whose image is finite. In other words Ω assigns to each state $s \in S$ a natural number, also referred to as a *priority*, taken from a finite pool of options $\{n_0, \dots, n_k\} = \Omega(S)$. We denote with $\max(\Omega)$, $\min(\Omega)$ and $|\Omega|$ the natural numbers $\max\{n_0, \dots, n_k\}$, $\min\{n_0, \dots, n_k\}$ and $|\{n_0, \dots, n_k\}|$ respectively.

The function Ω induces a set of completed paths, denoted by \mathcal{W}_Ω , specified as follows.

Definition 4.16. Let $\mathcal{A} = \langle (S, E), (S_1, S_2, S_N, B), \pi \rangle$ be a $2\frac{1}{2}$ -player tree game arena and Ω a parity assignment for it. A completed path \vec{s} belong to the *parity set induced by Ω* , denoted by $\mathcal{W}_\Omega \subseteq \mathcal{P}$, if either:

- (1) \vec{s} is a finite terminated path, i.e., $\vec{s} \in \mathcal{P}^t$, and the priority assigned to the last state of \vec{s} is odd, i.e., $\Omega(\text{last}(\vec{s})) \equiv 1 \pmod{2}$, or
- (2) \vec{s} is infinite, i.e., $\vec{s} \in \mathcal{P}^\omega$ with $\vec{s} = \{s_i\}_{i \in \mathbb{N}}$, and the greatest priority assigned to infinitely many states s_i in \vec{s} is even, i.e., $(\limsup_{i \in \mathbb{N}} \Omega(s_i)) \equiv 0 \pmod{2}$.

It is well known (see, e.g., [7]) that \mathcal{W}_Ω is a Δ_3^0 set, hence a Borel set.

Definition 4.17. A *player assignment* Pl for a two-player stochastic tree game arena $\mathcal{A} = \langle (S, E), (S_1, S_2, S_N, B), \pi \rangle$ is a function $\text{Pl} : B \rightarrow \{1, 2\}$. We often include the information provided by Pl directly in the signature of \mathcal{A} by considering the partition $(S_1, S_2, S_N, B_1, B_2)$, where $B_i = \text{Pl}^{-1}(\{i\})$, for $i \in \{1, 2\}$.

The function Pl assigns a *player identifier* to each state $s \in B$. This allows to consider each branching play T in \mathcal{A} as a parity game G_T (induced by Ω and Pl) played by Player 1 and Player 2 on the tree T , where Player 1 and Player 2 controls the vertices \vec{s} of T such that $\text{Pl}(\text{last}(\vec{s})) = 1$ and $\text{Pl}(\text{last}(\vec{s})) = 2$ respectively. All other vertices are either leaves, in which case the game G_T ends, or have a unique child, towards which the game automatically progresses. The result of the game is a branch in T or, equivalently, a completed path in

\mathcal{A} . Adopting standard terminology (see, e.g., [20], [30]), we say that Player 1 (respectively Player 2) wins the parity game G_T if they have a winning strategy guaranteeing the outcome of the game to be in the set \mathcal{W}_Ω (respectively, in $\mathcal{P} \setminus \mathcal{W}_\Omega$). It is well known (see, e.g., [20]), that one of the two player has a winning strategy in G_T .

We are now ready to formally introduce the notion of $2\frac{1}{2}$ -player meta-parity game.

Definition 4.18. Given a $2\frac{1}{2}$ -player tree game arena, a priority assignment Ω and a player assignment Pl for it, the associated *two-player stochastic meta-parity game* is defined as the $2\frac{1}{2}$ -player tree game $\mathcal{G} = \langle \mathcal{A}, \Phi_{\Omega, \text{Pl}} \rangle$, where $\Phi_{\Omega, \text{Pl}} \subseteq \mathcal{BP}$ is defined as follows:

$$\Phi_{\Omega, \text{Pl}} = \{T \mid T \in \mathcal{BP} \text{ and Player 1 has a winning strategy in } G_T\}.$$

Note that, by previous observations, the set $\mathcal{BP} \setminus \Phi_{\Omega, \text{Pl}}$ can be specified as follows:

$$\mathcal{BP} \setminus \Phi_{\Omega, \text{Pl}} = \{T \mid T \in \mathcal{BP} \text{ and Player 2 has a winning strategy in } G_T\}.$$

We often just write Φ if the priority and player assignments are clear from the context.

Remark 4.19. Note that a $2\frac{1}{2}$ -player meta-parity game $\mathcal{G} = \langle \mathcal{A}, \Phi_{\Omega, \text{Pl}} \rangle$ without branching states is just an ordinary $2\frac{1}{2}$ -player game. The game G_T associated with a branching play T (which is just a completed path \vec{s} , see Remark 4.14) is trivial, and belongs to $\Phi_{\Omega, \text{Pl}}$ if and only if $\vec{s} \in \mathcal{W}_\Omega$. Thus $2\frac{1}{2}$ -player meta-parity games generalize ordinary $2\frac{1}{2}$ -player parity games in an obvious way.

Note that Definition 4.18 is meaningful, in accordance with Definition 4.8, only if the set $\Phi_{\Omega, \text{Pl}}$ is a universally measurable set of branching plays. As we now discuss, this turns out to be a delicate point.

Theorem 4.20. *Given a $2\frac{1}{2}$ -player tree game arena \mathcal{A} , a priority assignment Ω and a player assignment Pl for it, the associated set $\Phi_{\Omega, \text{Pl}}$ of branching plays is a Δ_2^1 set.*

Proof. Let us denote with $\mathcal{P}_{B_i}^{<\omega}$ the set of finite paths $\vec{s} \in \mathcal{P}^{<\omega}$ in \mathcal{A} such that $\text{last}(\vec{s}) \in B$ and $\text{Pl}(\text{last}(\vec{s})) = i$, for $i \in \{1, 2\}$. Let us consider the set Σ_i of functions $\mathcal{P}_{B_i}^{<\omega} \rightarrow \mathcal{P}^{<\omega} \cup \{\bullet\}$. This set contains all the strategies available to Player i in every game \mathcal{G}_T , for $T \in \mathcal{BP}$, seen as functions $f \in \Sigma_i$ restricted to T . We endow Σ_i with the Baire space-like topology, where for every pair (x, y) , with $x \in \mathcal{P}_{B_i}^{<\omega}$ and $y \in \mathcal{P}^{<\omega} \cup \{\bullet\}$, the set $O_{x,y}$ of all functions $f \in \Sigma_i$ such that $f(x) = y$ is a basic open set. This is a 0-dimensional Polish space.

Let us now consider the subset of $\mathcal{BP} \times \Sigma_1 \times \Sigma_2$, denoted by \mathcal{T} , consisting of all triples (T, σ_1, σ_2) such that the strategies σ_1 and σ_2 are valid strategies in \mathcal{G}_T , i.e., functions from finite sequences of vertices in T to vertices in T . It is easy to see that \mathcal{T} is a closed subset of $\mathcal{BP} \times \Sigma_1 \times \Sigma_2$ (endowed with the product topology). Indeed, the set of triples which do not belong to \mathcal{T} is open, because one can tell if one of the two strategies σ_1 and σ_2 is not valid in the inner game \mathcal{G}_T , i.e., it makes choices which are not in T , just by looking at finite information about T , σ_1 and σ_2 . Hence \mathcal{T} is a Polish space, as it is a closed subset of the Polish space $\mathcal{BP} \times \Sigma_1 \times \Sigma_2$.

Let us denote with $\text{out} : \mathcal{T} \rightarrow \mathcal{P}$ the function which maps a triple (T, σ_1, σ_2) to the induced play (i.e., a completed path in \mathcal{A}) in \mathcal{G}_T . The function out is clearly continuous.

Let us now consider the set $A \subseteq \mathcal{T}$ defined as the set of triples (T, σ_1, σ_2) such that the game G_T is won by Player 2 when the two players follow the (valid for \mathcal{G}_T) strategies σ_1 and σ_2 respectively, i.e., the set formally defined as $A = \text{out}^{-1}(\mathcal{P} \setminus \mathcal{W}_\Omega)$. Since out is continuous and \mathcal{W}_Ω is a Borel set, it follows that A is a Borel set.

Let us now define the set $B \subseteq \mathcal{BP} \times \Sigma_1$ as $B = \{(T, \sigma_1) \mid \exists \sigma_2 \in \Sigma_2. (T, \sigma_1, \sigma_2) \in A\}$. The set B is the set of all pairs (T, σ_1) , such that Player 2 has a strategy σ_2 in the game \mathcal{G}_T

winning against σ_1 , i.e., such that the strategy profile (σ_1, σ_2) induces in \mathcal{G}_T a completed path in $\mathcal{P} \setminus \mathcal{W}_\Omega$. The set B is a Σ_1^1 set by construction. Observe that its complement \overline{B} is the set of all pairs (T, σ_1) such that Player 2 does not have a strategy σ_2 for the game \mathcal{G}_T winning against the strategy σ_1 . Equivalently, by determinacy of 2-player parity games [20], the strategy σ_1 is a winning strategy for Player 1 in the game \mathcal{G}_T . The set \overline{B} is a Π_1^1 set by construction. We can now define the set $\Phi_{\Omega, \text{Pl}} \subseteq \mathcal{BP}$ of all branching plays T where Player 1 has a winning strategy in \mathcal{G}_T as $\Phi_{\Omega, \text{Pl}} = \{T \mid \exists \sigma_1 \in \Sigma_1. (T, \sigma_1) \in \overline{B}\}$. It then follows that, by construction, $\Phi_{\Omega, \text{Pl}}$ is a Σ_2^1 set.

The desired result then follows by observing that the complement set $\overline{\Phi_{\Omega, \text{Pl}}}$ is also a Σ_2^1 set. This is because $\overline{\Phi_{\Omega, \text{Pl}}}$ is the winning set associated with the triplet $(\mathcal{A}, \overline{\Omega}, \overline{\text{Pl}})$, where $\overline{\Omega}$ is the dual parity assignment defined as $\overline{\Omega}(s) = \Omega(s) + 1$, and $\overline{\text{Pl}}$ is the dual player assignment inverting the role of the two players, specified as $\overline{\text{Pl}}(b) = 0$ if and only if $\text{Pl}(b) = 1$. \square

The following theorem asserts that the result of Theorem 4.20 is strict. Thus the technologies employed in this paper for dealing with the complexity of the winnings sets in $2\frac{1}{2}$ -player meta-parity games are not trivially avoidable. Since the technicalities required for proving the result would not be particularly useful for the main theorem of this paper, we just provide a reference to a detailed proof.

Theorem 4.21. *There exists a $2\frac{1}{2}$ -player meta-parity game $\mathcal{G} = \langle \mathcal{A}, \Phi_{\Omega, \text{Pl}} \rangle$, having a finite arena \mathcal{A} , such that the winning set $\Phi_{\Omega, \text{Pl}}$ is not analytic nor co-analytic, i.e., $\Phi_{\Omega, \text{Pl}} \notin \Sigma_1^1 \cup \Pi_1^1$.*

Proof. One can construct an explicit example of finite $2\frac{1}{2}$ -player meta-parity game and show that both a Σ_1^1 -complete set and a Π_1^1 -complete set are Wadge-reducible to the winning set $\Phi_{\Omega, \text{Pl}}$. We refer to Theorem 6.4.3 in [25] for a detailed proof. \square

It follows that, absent any further evidence, the winning set of a $2\frac{1}{2}$ -player meta-parity game might be not provably universally measurable in ZFC set theory. However, as stated in Theorem 2.10, it is provably universally measurable in $\text{ZFC} + \text{MA}_{\aleph_1}$ set theory. This is one of the uses we make of Martin's axiom at \aleph_1 but, as we shall note later in Section 8.1, not the only one.

The following theorem, which will be useful later, exposes an important property of $2\frac{1}{2}$ -player meta-parity games and sheds some light on the relationship between this class of games and the logic $\text{pL}\mu^\odot$. It constitutes the expected generalization of the corresponding property of ordinary $2\frac{1}{2}$ -player parity games (see, e.g., Proposition 4.19 in [26]).

Theorem 4.22. *Let $\mathcal{G} = \langle \mathcal{A}, \Phi_{\Omega, \text{Pl}} \rangle$ be a two-player stochastic parity game with arena $\mathcal{A} = \langle (S, E), (S_1, S_2, S_N, B_1, B_2), \pi \rangle$, where $B_i = \text{Pl}^{-1}(\{i\})$. The functions $\text{VAL}_\downarrow(\mathcal{G})$ and $\text{VAL}_\uparrow(\mathcal{G})$, of type $S \rightarrow [0, 1]$, are fixed points of the functional $\mathcal{F}: [0, 1]^S \rightarrow [0, 1]^S$ defined as follows:*

$$\mathcal{F}(f)(s) = \begin{cases} \Omega(s) \pmod{2} & \text{if } E(s) = \emptyset, \text{ i.e., if } s \text{ is a terminal state} \\ \bigsqcup_{t \in E(s)} f(t) & \text{if } s \in S_1 \\ \prod_{t \in E(s)} f(t) & \text{if } s \in S_2 \\ \sum_{t \in E(s)} \pi(s)(t) \cdot f(t) & \text{if } s \in S_N \\ \prod_{t \in E(s)} f(t) & \text{if } s \in B_1 \\ \prod_{t \in E(s)} f(t) & \text{if } s \in B_2 \end{cases}$$


Proof. The proof is carried out following the same methodology of, e.g., Proposition 4.19 in [26]. The interesting cases are associated with the analysis of the branching nodes $s \in B_1$ and $s \in B_2$. In what follows we just show that, for every $s \in B_1$, the following equality holds:

$$\text{VAL}_\downarrow(\mathcal{G})(s) \geq \prod_{t \in E(s)} \text{VAL}_\downarrow(\mathcal{G})(t).$$

The reverse inequality $\text{VAL}_\downarrow(\mathcal{G})(s) \leq \prod_{t \in E(s)} \text{VAL}_\downarrow(\mathcal{G})(t)$, as well as all other cases for $s \in \{S_1, S_2, S_N, B_2\}$ can be proved in a similar way. We refer to Theorem 5.2.10 of [25] for a detailed proof covering all cases.

By Definition 4.12 of $\text{VAL}_\downarrow(\mathcal{G})$, we need to prove that the equality $\bigsqcup_{\sigma_1} \prod_{\sigma_2} \mathbf{E}(M_{\sigma_1, \sigma_2}^s) \geq \prod_{t \in E(s)} (\bigsqcup_{\tau_1} \prod_{\tau_2} \mathbf{E}(M_{\tau_1, \tau_2}^t))$ holds. Let $E(s) = \{t_i\}_{i \in I}$. At the state s the game is split in I -many subplays continuing their execution from the states t_i . Let τ_1^i be a ε_i -optimal strategy (see Definition 4.13) for Player 1, with $\varepsilon_i > 0$.

Define the strategy σ_1 for Player 1, when the game starts at s and the I -many subplays are generated, to behave in the subplay continuing its execution from t_i as the strategy τ_1^i . Given any strategy σ_2 for Player 2, the Markov branching play M_{σ_1, σ_2}^s can be depicted² as follows:




where, for each $i \in I$, the strategy τ_2^i is specified as $\tau_2^i(\vec{t}_i) = \sigma_2(s, \vec{t}_i)$ for all finite paths \vec{t}_i with $\text{first}(\vec{t}_i) = t_i$. Let us denote, to improve readability, with M and M_i , for $i \in I$, the Markov branching plays M_{σ_1, σ_2}^s and $M_{\tau_1^i, \tau_2^i}^{t_i}$, respectively. We are now going to show that the following equality holds:

$$\mathbb{P}_M(\Phi_{\Omega, \mathbf{P}1}) = \prod_{i \in I} \mathbb{P}_{M_i}(\Phi_{\Omega, \mathbf{P}1}) \quad (4.1)$$

²The edge connecting s with t_i has not been dashed to highlight that s is a branching node.

holds. This will conclude the proof. Indeed note that, by construction of σ_1^i , the inequality $\mathbb{P}_{M_i}(\Phi_{\Omega, \text{Pl}}) \geq \text{VAL}_{\downarrow}(\mathcal{G})(t_i) - \varepsilon_i$ holds. Thus, the strategy σ_1 guarantees, by appropriate choices of values ε_i , for $i \in I$, a value closed to $\prod_{i \in I} \text{VAL}_{\downarrow}(\mathcal{G})(t_i)$ as desired.

For what concerns \mathbb{P}_M , we can restrict attention to the set of branching plays \mathcal{BP}_s , where \mathcal{BP}_s denotes the set of branching plays in \mathcal{A} rooted at s , since the set of all other branching plays in \mathcal{G} gets assigned probability 0 by Definition 4.7 of \mathbb{P}_M . Similarly, when considering \mathbb{P}_{M_i} , we can restrict to the set \mathcal{BP}_i of branching plays rooted at t_i . We can depict the branching plays in \mathcal{BP}_s and the branching plays in \mathcal{BP}_i as follows:



where we use T_i to range over the set of branching plays in \mathcal{G} rooted at t_i . We denote with $s[T_i]_{i \in I}$ the branching play on the left. Let $\prod_i \mathcal{BP}_i$ be endowed with the product topology. Define $m : \prod_i \mathcal{BP}_i \rightarrow \mathcal{BP}_s$ as $m(\{T_i\}_{i \in I}) = s[T_i]_{i \in I}$. It is easy to verify that m is a homeomorphism. Consider the product measure $\times_{i \in I} \mathbb{P}_{M_i}$ on the space $\prod_i \mathcal{BP}_i$. We now show that the equality

$$\times_{i \in I} \mathbb{P}_{M_i}(X) = \mathbb{P}_M(m(X)) \quad (4.2)$$

holds for every measurable $X \subseteq \prod_i \mathcal{BP}_i$. By regularity of measures in Polish spaces (see, e.g., [17]), we just need to prove that for each basic open set $O \subseteq \prod_i \mathcal{BP}_i$ the equality $\times_{i \in I} \mathbb{P}_{M_i}(O) = \mathbb{P}_M(m(O))$ holds. The basic sets O in the product topology are of the form $O_{F_0} \times \dots \times O_{F_k} \times \prod_{i > k} \mathcal{BP}_i$ with $O_{F_n} \subseteq \mathcal{BP}_n$, for some $k \in \mathbb{N}$ and $0 \leq n \leq k$. As usual, O_{F_n} denotes the basic open set of branching plays containing the finite tree F_n . Equality 4.2 then follows by definition of m and Definition 4.7 of the probability measures \mathbb{P}_M and \mathbb{P}_{M_i} .

The validity of Equation 4.1 is then derived as follows:

$$\begin{aligned} \mathbb{P}_M(\Phi_{\Omega, \text{Pl}}) &=_{\text{A}} \mathbb{P}_M(\Phi_{\Omega, \text{Pl}} \cap \mathcal{BP}_s) \\ &=_{\text{B}} \mathbb{P}_M(m(\overline{\prod_{i \in I} (\Phi_{\Omega, \text{Pl}} \cap \mathcal{BP}_i)})) \\ &=_{\text{C}} 1 - \mathbb{P}_M(m(\prod_{i \in I} (\overline{\Phi_{\Omega, \text{Pl}} \cap \mathcal{BP}_i}))) \\ &=_{\text{D}} 1 - \times_{i \in I} \mathbb{P}_{M_i}(\overline{\Phi_{\Omega, \text{Pl}} \cap \mathcal{BP}_i}) \\ &=_{\text{E}} 1 - \prod_{i \in I} (1 - \mathbb{P}_{M_i}(\Phi_{\Omega, \text{Pl}} \cap \mathcal{BP}_i)) \\ &=_{\text{F}} 1 - \prod_{i \in I} (1 - \mathbb{P}_{M_i}(\Phi_{\Omega, \text{Pl}})) \\ &=_{\text{G}} \prod_{i \in I} \mathbb{P}_{M_i}(\Phi_{\Omega, \text{Pl}}). \end{aligned}$$

where the overlined sets denote the expected complements. Steps (A) and (F) hold by previous observations. To justify step (B), observe (using a standard strategy stealing argument) that a branching play $s[T_i]_{i \in I}$ is winning for Player 1 (i.e., it is in $\Phi_{\Omega, \text{Pl}}$) if and only if at least one of its subtrees T_i is winning for Player 1. This is because the state $s \in B_1$ is under the control of Player 1 in the inner game. Step (C) is valid because m is a homeomorphism, thus a bijection. Step (D) holds by Equation 4.2. Step (E) holds by definition of product measure. Lastly, step (G) holds by De Morgan dualities of the operations of product and coproduct. \square

We conclude this section by remarking that the notion of $2\frac{1}{2}$ -player meta-parity game can be further generalized, allowing the inner games to be 2-player games with general Borel winning sets. We refer to [25, §5] for an analysis of this sort of games.

5. GAME SEMANTICS OF $\text{pL}\mu^\odot$

In this section we define the game semantics of the logic $\text{pL}\mu^\odot$. As for $\text{L}\mu$ [30] and $\text{pL}\mu$ [23, 26], given a PLTS $\langle P, \{\xrightarrow{a}\}_{a \in L} \rangle$ and an interpretation of the variables ρ a game \mathcal{G}_ρ^F is constructed for every $\text{pL}\mu^\odot$ formula F . The game semantics of a formula F at a state p is defined as the value of the game at a designated state $\langle p, F \rangle$. The logics $\text{L}\mu$ and $\text{pL}\mu$ are interpreted using ordinary 2-player parity games and $2\frac{1}{2}$ -player parity games, respectively. As anticipated earlier, we shall interpret $\text{pL}\mu^\odot$ with the novel class of $2\frac{1}{2}$ -player meta-parity games. Following the approach of [30], we first identify a class of $\text{pL}\mu^\odot$ formulas which is easier to work with and allow the simplification of some definitions.

Definition 5.1 ([30]). We say that a $\text{pL}\mu^\odot$ formula F is in *normal form*, if every occurrence of a μ or ν binder binds a distinct variable, and no variable appears both free and bound. Every formula can be put in normal form by standard α -renaming of the bound variables.

Definition 5.2 ([30]). Given a $\text{pL}\mu^\odot$ formula F in normal form, we say that a variable X *subsumes* a variable Y if X and Y are bound in F by the subformulas $\star_1 X.G$ and $\star_2 Y.H$ respectively, and $\star_2 Y.H \in \text{Sub}(G)$, for $\star_1, \star_2 \in \{\mu, \nu\}$.

Definition 5.3. We say that a $\text{pL}\mu^\odot$ formula F is in *product normal form* if

- (1) it is in normal form, and
- (2) it does not contain subformulas of the form $G \cdot G$ or $G \odot G$.

Every formula can be put in product normal form by (recursively) replacing every subformula $G \star G$ with, e.g., the semantically equivalent formula $G \star (G \vee G)$, for $\star \in \{\cdot, \odot\}$.

In what follows we shall restrict our attention, without loss of generality, to formulas in product normal form. While the advantage of restricting to normal formulas is clear [30], the product normal form is useful for the following reason. As anticipated in the introduction, we shall interpret game-states of the form $\langle p, F \cdot G \rangle$ and $\langle p, F \odot G \rangle$ as branching states having the states $\langle p, F \rangle$ and $\langle p, G \rangle$ as children. When $F = G$, the set of children becomes a singleton, and this does not reflect the intended game interpretation. The product normal form is one of the simplest way to avoid this kind of situations. As an alternative solution, one might consider *multisets* of successor states, rather than sets, in the definition of $2\frac{1}{2}$ -player meta-parity games.

We are now ready to specify how, given a PLTS $\mathcal{L} = \langle P, \{\xrightarrow{a}\}_{a \in L} \rangle$, the $2\frac{1}{2}$ -player meta-parity game $\mathcal{G}_\rho^F = \langle \mathcal{A}, \Phi_{\Omega, \text{Pl}} \rangle$, associated with a $\text{pL}\mu^\odot$ formula F and a $[0, 1]$ -interpretation ρ of the variables, is constructed.

The game arena $\mathcal{A} = \langle (S, E), (S_1, S_2, S_N, B), \pi \rangle$ is defined as follows. The set of game states S is defined as $S \stackrel{\text{def}}{=} (P \times \text{Sub}(F)) \cup (\mathcal{D}(\mathcal{L}) \times \text{Sub}(F)) \cup \{\top, \perp\}$, where $\mathcal{D}(\mathcal{L})$ denotes the (countable) set of probability distribution in \mathcal{L} (see Definition 2.2), and $\{\top, \perp\}$ are two special states representing immediate win and loss for Player 1, respectively. The relation E is defined as follows: $E(\langle d, G \rangle) \stackrel{\text{def}}{=} \{\langle p, G \rangle \mid p \in \text{supp}(d)\}$ for every $d \in \mathcal{D}(\mathcal{L})$. The set $E(\langle p, G \rangle)$ is defined, by case analysis on the outermost connective of G , as follows:

- (1) if $G = X$, with X free in F , then $E(\langle p, G \rangle) \stackrel{\text{def}}{=} \{\perp, \top\}$.
- (2) if $G = X$, with X bound in F by the subformula $\star X.H$, with $\star \in \{\mu, \nu\}$, then $E(\langle p, G \rangle) \stackrel{\text{def}}{=} \{\langle p, H \rangle\}$.
- (3) if $G = \star X.H$, with $\star \in \{\mu, \nu\}$, then $E(\langle p, G \rangle) \stackrel{\text{def}}{=} \{\langle p, H \rangle\}$.
- (4) if $G = \langle a \rangle H$, then $E(\langle p, G \rangle) \stackrel{\text{def}}{=} \{\langle d, H \rangle \mid p \xrightarrow{a} d\}$.
- (5) if $G = [a] H$, then $E(\langle p, G \rangle) \stackrel{\text{def}}{=} \{\langle d, H \rangle \mid p \xrightarrow{a} d\}$.
- (6) if $G = H \star H'$ with $\star \in \{\vee, \wedge, \odot, \cdot\}$ then $E(\langle p, G \rangle) = \{\langle p, H \rangle, \langle p, H' \rangle\}$

Lastly we define \top and \perp to be terminal states, i.e., $E(\perp) = E(\top) = \emptyset$. The partition (S_1, S_2, S_N, B) of S is defined as follows: every state $\langle p, G \rangle$ with G 's main connective in $\{\langle a \rangle, \vee, \mu X\}$ or with $G = X$ where X is a μ -variable, is in S_1 . Dually, every state $\langle p, G \rangle$ with G 's main connective in $\{[a], \wedge, \nu X\}$ or with $G = X$ where X is a ν -variable, is in S_2 . Every state of the form $\langle d, G \rangle$ or $\langle p, X \rangle$, with X free in F , is in S_N . Every state $\langle p, G \rangle$ whose G 's main connective is either \cdot or \odot is in B . Lastly we define the terminal states \perp and \top to be in S_1 and S_2 respectively. The function $\pi: S_N \rightarrow \mathcal{D}(S)$ assigns a probability distribution to every state under the control of Nature (thus specifying its intended probabilistic behavior) and it is defined as $\pi(\langle d, G \rangle)(\langle p, G \rangle) = d(p)$ on all states of the form $\langle d, G \rangle$. All other states in S_N are of the form $\langle p, X \rangle$, with X free in F . The function π is defined on these states as follows:

$$\pi(\langle p, X \rangle)(s) \stackrel{\text{def}}{=} \begin{cases} \rho(X)(p) & \text{if } s = \top \\ 1 - \rho(X)(p) & \text{if } s = \perp \\ 0 & \text{otherwise} \end{cases}$$

The priority assignment $\Omega: S \rightarrow \mathbb{N}$ is defined as usual in μ -calculus games: an odd priority is assigned to the states $\langle p, X \rangle$ with X a μ -variable and, dually, an even priority is assigned to the states $\langle p, X \rangle$ with X a ν -variable, in such a way that if Z subsumes Y in F then $\Omega(\langle p, Z \rangle) > \Omega(\langle p, Y \rangle)$. Moreover, for every terminal state $s \in S$, we define $\Omega(s) = 0$ if $s \in S_1$, and $\Omega(s) = 1$ if $s \in S_2$. This implements the policy that a player who gets stuck at a terminal state loses (see Theorem 4.22). All other states get assigned, by convention, priority 0. Lastly, the player assignment $\text{Pl}: B \rightarrow \{1, 2\}$ is defined as $\text{Pl}(\langle p, G_1 \odot G_2 \rangle) = 1$ and $\text{Pl}(\langle p, G_1 \cdot G_2 \rangle) = 2$, for every $p \in P$ and $G_1, G_2 \in \text{Sub}(F)$.

Remark 5.4. We now list some useful observations about the above defined game \mathcal{G}_ρ^F .

- (1) If no (co)product operators occur in F , then $B = \emptyset$, and the game \mathcal{G}_ρ^F is equivalent to the $\text{pL}\mu$ games of [23, 26] (see Remark 4.19). Thus the game semantics for $\text{pL}\mu^\odot$ generalizes the game semantics of $\text{pL}\mu$, as previously claimed.
- (2) If the PLTS \mathcal{L} is finite then the game arena \mathcal{A} of \mathcal{G}_ρ^F is finite too.
- (3) Note how the only role of two game states $\{\top, \perp\}$ is to interpret the game states $\langle p, X \rangle$ with X a free variable. By application of Theorem 4.22, it follows that the game values at \top , \perp and $\langle p, X \rangle$ are 1, 0, and $\rho(X)(p)$ respectively, as expected. This solution avoids the otherwise necessary introduction of $[0, 1]$ -valued payoff functions in the definition of $2\frac{1}{2}$ -player meta-parity games.


We are now ready to state our main result.

Theorem 5.5 (MA_{\aleph_1}). *Given a PLTS $\mathcal{L} = \langle P, \{\xrightarrow{a}\}_{a \in L} \rangle$, for every state $p \in P$, interpretation of the variables ρ and $\text{pL}\mu^\odot$ formula F , the following equalities hold:*


$$\text{VAL}\downarrow(\mathcal{G}_\rho^F)(\langle p, F \rangle) = \text{VAL}\uparrow(\mathcal{G}_\rho^F)(\langle p, F \rangle) = \llbracket F \rrbracket_\rho(p).$$

Note that Theorem 5.5 asserts that $\text{pL}\mu^\odot$ games are determined. In light of this result, we can just write $\text{VAL}(\mathcal{G}_\rho^F)$ for the unique value function of the game \mathcal{G}_ρ^F . The *game semantics* of F at p under the interpretation ρ is then defined as $\text{VAL}(\mathcal{G}_\rho^F)(\langle p, F \rangle)$, and it coincides with the denotational semantics $\llbracket F \rrbracket_\rho(p)$, as desired.

The game semantics for $\text{pL}\mu^\odot$ offers a clear and simple operational interpretation for the meaning of the qualitative threshold modalities $\mathbb{P}_{>0}$ and $\mathbb{P}_{=1}$ (see Definition 3.3). Let us consider, for example, the game $\mathcal{G}_\rho^{\mathbb{P}_{>0}F}$ associated with a PLTS $\mathcal{L} = \langle P, \{\xrightarrow{a}\}_{a \in L} \rangle$, a $[0, 1]$ -interpretation ρ and the $\text{pL}\mu^\odot$ formula $\mathbb{P}_{>0}F \stackrel{\text{def}}{=} \mu X.(F \odot X)$. The game $\mathcal{G}_\rho^{\mathbb{P}_{>0}F}$, at the state $\langle p, \mu X.(F \odot X) \rangle$, for some $p \in P$, can be depicted as follows:



After an initial unfolding step from $\langle p, \mu X.(F \odot X) \rangle$ to $\langle p, F \odot X \rangle$, the game is split in two concurrent sub-games, one continuing its execution from the state $\langle p, F \rangle$ (this sub-game can be considered an instance of the game \mathcal{G}_ρ^F starting at $\langle p, F \rangle$) and the other from the state $\langle p, X \rangle$. In order to win the game $\mathcal{G}_\rho^{\mathbb{P}_{>0}F}$, Player 1 has to win in at least one of the two generated sub-games, thus either in the instance of \mathcal{G}_ρ^F or in the sub-games continuing at $\langle p, X \rangle$. This second sub-game, however, after an unfolding step, progresses to the game state $\langle p, F \odot X \rangle$, where the protocol is repeated generating yet another two sub-games. The infinite execution of the game leads to the generation of infinitely many instances of the game \mathcal{G}_ρ^F . A branching play T in $\mathcal{G}_\rho^{\mathbb{P}_{>0}F}$ can be depicted as follows:



where T_1, T_2, \dots , represent the branching plays corresponding to the plays in each generated instance of the game \mathcal{G}_ρ^F . Since the variable X unfolded infinitely often in the rightmost path in T is bound by a least fixed point in $\mu X.(F \odot X)$, and since the \odot nodes are Player 1 choices in the inner game \mathcal{G}_T , we have that T is a winning branching play for Player 1 if and only if there exists some $n \in \mathbb{N}$ such that T_n , the outcome of the n -th generated instance of $\mathcal{G}(F, \rho)$, is winning for Player 1. Thus, the game $\mathcal{G}_\rho^{\mathbb{P}_{>0}F}$, at the state $\langle p, \mu X.F \odot X \rangle$ can be simply described as follows: generate an infinite number of instances of the game \mathcal{G}_ρ^F at the state $\langle p, F \rangle$; Player 1 wins if at least one of the infinitely many generated instances of

\mathcal{G}_ρ^F ends up in a winning branching play and Player 2 wins otherwise. It is then quite clear that if $\text{VAL}(\mathcal{G}_\rho^{\mathbb{P}_{>0}F})(\langle p, \mathbb{P}_{>0}F \rangle) > 0$ (or, equivalently, $\llbracket F \rrbracket_\rho(p) > 0$ by Theorem 5.5), then the probability that at least one (and in fact countably many) of the infinitely many instances of \mathcal{G}_ρ^F will result in a win for Player 1, is 1. Similarly, if $\llbracket F \rrbracket_\rho(p) = 0$, then the probability that at least one of the instances of \mathcal{G}_ρ^F will result in a win for Player 1, is 0.

The game semantics for $\text{pL}\mu^\odot$ thus offers a straightforward interpretation for the probabilistic qualitative modality $\mathbb{P}_{>0}$ exploiting the simple idea that an event (which we can, at some extent, see as a $\text{pL}\mu^\odot$ property) has probability greater than zero if and only if, when repeated infinitely many time, it almost surely occurs at least once. An analogous straightforward interpretation can be given to the other qualitative threshold modality $\mathbb{P}_{=1}F \stackrel{\text{def}}{=} \nu X.(F \cdot X)$: generate an infinite number of instances of the game \mathcal{G}_ρ^F at the state $\langle p, F \rangle$; Player 1 wins if all of them end up in a winning branching play for Player 1, and Player 2 wins otherwise.

6. INDUCTIVE CHARACTERIZATION OF THE WINNING SET OF $\mathcal{G}_\rho^{\mu X.F}$

In this section we provide a transfinite inductive characterization of the set $\Phi^{\mu X.F}$ of winning branching plays of the game $\mathcal{G}_\rho^{\mu X.F}$, for an arbitrary $\text{pL}\mu^\odot$ formula F and interpretation ρ . This result will be used in the proof of our main result, in Section 8. The inductive characterization of $\Phi^{\mu X.F}$ will be obtained exploiting the similarities between the $\text{pL}\mu^\odot$ games $\mathcal{G}_\rho^{\mu X.F}$ and the simpler (in the complexity of the formula) game \mathcal{G}_ρ^F . The game arenas of the two games, denoted here by $\mathcal{A}^{\mu X.F}$ and \mathcal{A}^F , are almost identical as they differ only in the following aspects:

- (1) The set of game states of $\mathcal{A}^{\mu X.F}$ is the set of states of \mathcal{A}^F plus the set of states of the form $\langle p, \mu X.F \rangle$. The latter states, however, play almost no role in the game because they have only one successor, namely $\langle p, F \rangle$, and are not reachable by any other state.
- (2) More importantly, the states of the form $\langle p, X \rangle$, which are present in both game arenas, are Player 1 states in $\mathcal{G}_\rho^{\mu X.F}$ (with $\langle p, F \rangle$ as unique successor), and Nature states in \mathcal{G}_ρ^F (with two³ terminal successors, \top and \perp , reachable with probability $\rho(X)(p)$ and $1 - \rho(X)(p)$, respectively).

Moreover observe that the player assignments of the two games, denoted here by Pl^F and $\text{Pl}^{\mu X.F}$, are identical, and that the priority assignments, Ω^F and $\Omega^{\mu X.F}$, differ only on the game-states s of the form $\langle p, X \rangle$: $\Omega^F(s) = 0$ and $\Omega^{\mu X.F}(s) = p$ for some odd (maximal) priority $p = \max(\Omega^{\mu X.F})$ (see Definition 4.15).

A branching play T in the game \mathcal{G}_ρ^F , rooted at the game state $\langle p, F \rangle$, can be depicted as in figure 3(a). The triangle represents⁴ the set of paths in T never reaching a state of the form $\langle q, X \rangle$, for $q \in P$, and the other edges represents the (possibly empty) collection of paths $\{\tilde{s}_i\}_{i \in I}$ in T , for some countable index set I , reaching a state of the form $\langle p_i, X \rangle$ which is (necessarily) followed by a terminal state $b_i \in \{\top, \perp\}$. Similarly a branching play

³Note that the set of successor state of $\langle p, X \rangle$ is defined to be $\{\top, \perp\}$, even when $\rho(X)(p) \in \{0, 1\}$.

⁴The picture is quite simplistic. For example there are, in general, paths in T never reaching a state of the form $\langle q, X \rangle$, that nonetheless branch away from the paths \tilde{s}_i , for $i \in I$, somewhere between the root $\langle p, F \rangle$ and the last state $\langle p_i, X \rangle$, whereas the picture depicts all such branches as branching away immediately at the root $\langle p, F \rangle$.

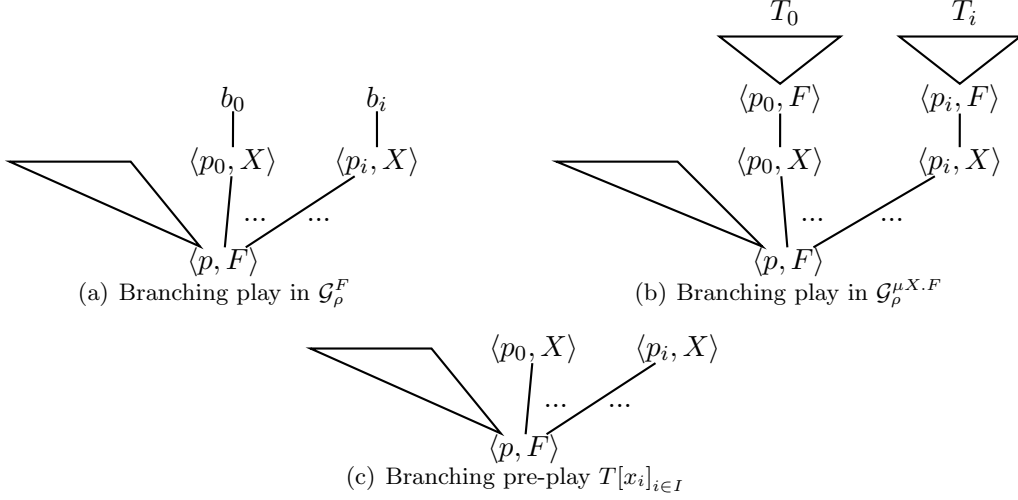


Figure 3: Branching plays and pre-plays

T in $\mathcal{G}_\rho^{\mu X.F}$, rooted at $\langle p, F \rangle$, can be depicted as in figure 3(b). We extract the common part between the branching plays of \mathcal{G}_ρ^F and $\mathcal{G}_\rho^{\mu X.F}$ by the following definition.

Definition 6.1 (Branching pre-play). Let T be a branching play in \mathcal{G}_ρ^F or in $\mathcal{G}_\rho^{\mu X.F}$, and $\{\vec{s}_i\}_{i \in I}$ be the I -indexed set of paths in T reaching a state of the form $\langle p_i, X \rangle$, as described above. The *branching pre-play* $T[x_i]_{i \in I}$, which can be depicted as in figure 3(c), is the tree obtained from the branching play T by pruning its subtrees rooted at \vec{s}_i , for $i \in I$.

The notation adopted for branching pre-plays is motivated by the use we make of them. We consider a branching pre-play $T[x_i]_{i \in I}$ as a context on which we can plug in other branching plays: at the hole indexed by i , for $i \in I$, any branching play rooted at $\langle p_i, X \rangle$ can be plugged. If we desire to obtain a branching play for the game \mathcal{G}_ρ^F , we shall fill the branching pre-play $T[x_i]_{i \in I}$ with branching plays T_i rooted at $\langle p_i, X \rangle$ in \mathcal{A}^F . These are trees ending immediately either in the leaf \top or in the leaf \perp . We denote with $T[b_i]_{i \in I}$ the branching play in \mathcal{G}_ρ^F obtained by filling the hole indexed by i with the branching play rooted at $\langle p_i, X \rangle$ and having b_i as leaf, for $b_i \in \{\top, \perp\}$ (see Figure 3(a)). Similarly given a I -indexed family $\{T_i\}_{i \in I}$ of branching plays in $\mathcal{G}_\rho^{\mu X.F}$, where each T_i is rooted in $\langle p_i, X \rangle$, we denote with $T[T_i]_{i \in I}$ the branching play in $\mathcal{G}_\rho^{\mu X.F}$ obtained by filling the hole indexed by i with the branching play T_i (see Figure 3(b)).

Clearly every branching play T in \mathcal{G}_ρ^F , rooted at $\langle p, F \rangle$, is uniquely of the form $T[b_i]_{i \in I}$ for an appropriate sequence $\{b_i\}_{i \in I}$. Similarly every branching play T rooted at $\langle p, F \rangle$ in $\mathcal{G}_\rho^{\mu X.F}$ is of the form $T[T_i]_{i \in I}$ for an appropriate sequence $\{T_i\}_{i \in I}$. We now exploit these observations to define the following function from branching plays in $\mathcal{G}_\rho^{\mu X.F}$ to branching plays in \mathcal{G}_ρ^F .

Definition 6.2. Let $X \subseteq \mathcal{BP}^{\mu X.F}$ be a set of branching plays in the game $\mathcal{G}_\rho^{\mu X.F}$. We define the function $m_X : \mathcal{BP}^{\mu X.F} \rightarrow \mathcal{BP}^F$, from branching plays in $\mathcal{G}_\rho^{\mu X.F}$ to branching plays in \mathcal{G}_ρ^F , as follows:

$$m_X(T[T_i]_{i \in I}) = T[T_i \in X]_{i \in I} \quad \text{where } T_i \in X \stackrel{\text{def}}{=} \begin{cases} \top & \text{if } T_i \in X \\ \perp & \text{otherwise} \end{cases}$$

Thus, the function m_X maps a branching play (uniquely expressible as $T[T_i]_{i \in I}$, as observed above) to the corresponding branching play $T[b_i]_{i \in I}$ in \mathcal{G}_ρ^F , obtained by filling the i -th hole of $T[x_i]_{i \in I}$, with the branching play having \top as leaf if and only if T_i belongs to the set X .

Lemma 6.3. *Let $X \subseteq \mathcal{BP}^{\mu X.F}$ be an open (Borel, universally measurable) set of branching plays in the game $\mathcal{G}_\rho^{\mu X.F}$. Then the function m_X is continuous (Borel measurable, universally measurable).*

Proof. It is clear that the function $\lambda T.(T \in X)$ of Definition 6.2 is continuous (Borel, universally measurable) precisely when X is open (Borel, universally measurable), where $\{\top, \perp\}$ is endowed with the discrete topology. The proof is then carried out with standard arguments (see, e.g., Theorem 6.2.18 in [25]). Here we omit the routine details. \square

We are now ready to define the operator of which $\Phi^{\mu X.F}$, the winning set of the game $\mathcal{G}_\rho^{\mu X.F}$, is the least fixed point.

Definition 6.4. The operator $\mathbb{W}: \mathcal{BP}^{\mu X.F} \rightarrow \mathcal{BP}^{\mu X.F}$ is defined as follows:

$$\mathbb{W}(X) \stackrel{\text{def}}{=} m_X^{-1}(\Phi^F) = \{T[T_i]_{i \in I} \mid T[T_i \in X]_{i \in I} \in \Phi^F\}$$

where Φ^F is the winning set of the game \mathcal{G}_ρ^F .

We now prove a few important properties of the operator \mathbb{W} .

Lemma 6.5. *If X is a set of branching plays winning for Player 1 in $\mathcal{G}_\rho^{\mu X.F}$, i.e., if $X \subseteq \Phi^{\mu X.F}$, then $\mathbb{W}(X) \subseteq \Phi^{\mu X.F}$ too.*

Proof. Fix some $X \subseteq \Phi^{\mu X.F}$ and consider an arbitrary branching play $T_1 = T[T_i]_{i \in I} \in \mathbb{W}(X)$. It follows by definition of \mathbb{W} that $T_2 = T[T_i \in X]_{i \in I} \in \Phi^F$, i.e., T_2 is a winning branching play in the game \mathcal{G}_ρ^F . Equivalently, Player 1 has a winning strategy σ in the inner game G_{T_2} (see Section 4.1). We now prove that $T_1 \in \Phi^{\mu X.F}$, i.e., that Player 1 has a winning strategy τ in the inner game G_{T_1} (see Section 4.1) by a strategy stealing argument, exploiting the common structure (the branching pre-play $T[x_i]_{i \in I}$) of the two branching plays T_1 and T_2 . The strategy τ behaves as the strategy σ until a hole \vec{s}_i , for $i \in I$, is reached. Thus if no hole is ever reached, the plays in the two games are identical, hence Player 1 wins following τ , as desired. If a hole $\vec{s}_i \in I$ is reached, then it is necessarily the case that $T_i \in X$, because otherwise the play in G_{T_2} would end in the losing state $\perp = T_i \in X$ while playing in accordance with the winning strategy σ . A contradiction. We then define the strategy τ to play the rest of the game as a winning strategy τ_i for the inner game \mathcal{G}_{T_i} . Note that such a strategy exist because $T_i \in X \subseteq \Phi^{\mu X.F}$. It then follows that Player 1, playing in accordance with τ , always produces a play with a winning tail. Thus, τ is winning, as desired. \square

Lemma 6.6. *The operator \mathbb{W} is monotone, i.e., $\mathbb{W}(X) \subseteq \mathbb{W}(Y)$ holds for every $X \subseteq Y$.*

Proof. Fix $X \subseteq Y \subseteq \mathcal{BP}^{\mu X.F}$. Assume $T_1 = T[T_i]_{i \in I} \in \mathbb{W}(X)$, i.e., $T[T_i \in X]_{i \in I} \in \Phi^F$. We need to prove that $T_1 \in \mathbb{W}(Y)$ too, i.e., that $T_2 = T[T_i \in Y]_{i \in I} \in \Phi^F$. Equivalently, we need to show that if Player 1 has a winning strategy in the inner game G_{T_1} associated with T_1 , then they have a winning strategy in G_{T_2} too. The two parity games G_{T_1} and G_{T_2} are almost identical, except that a play ending in one of the holes \vec{s}_i , for $i \in I$, might be losing for Player 1 in G_{T_1} (when the game ends in the leaf $\perp = T_i \in X$) but winning in G_{T_2} (when $T_i \in Y = \top$). The desired result then trivially follows. \square

As a consequence, by application of the Knaster–Tarski theorem, the operator \mathbb{W} has a least fixed point $\text{lfp}(\mathbb{W})$. We are now ready to prove the main result of this section.

Theorem 6.7. *The following equality holds: $\Phi^{\mu X.F} = \text{lfp}(\mathbb{W})$.*

Proof. We already know, by application of Lemma 6.5, that $\text{lfp}(\mathbb{W}) \subseteq \Phi^{\mu X.F}$. We now prove that the equality holds by showing that, for every $T \notin \text{lfp}(\mathbb{W})$, the branching play T does not belong to $\Phi^{\mu X.F}$. Fix an arbitrary $T \notin \text{lfp}(\mathbb{W})$. We show that $T \notin \Phi^{\mu X.F}$ by constructing a winning strategy σ_2^T for Player 2 in the inner game G_T . By assumption we have that $T = T[T_i]_{i \in I} \notin \text{lfp}(\mathbb{W})$ or, equivalently, $T[T_i \in \text{lfp}(\mathbb{W})]_{i \in I} \notin \Phi^F$. For notational convenience, let us denote with R the branching play $T[T_i \in \text{lfp}(\mathbb{W})]_{i \in I}$ in \mathcal{G}_ρ^F . Let τ_2^R be a strategy winning for Player 2 in the inner game G_R . As already done in the proof of Lemma 6.5, we define σ_2^T exploiting the common structure (the branching pre-play $T[x_i]_{i \in I}$) of T and R . The strategy σ_2^T behaves as the strategy τ_2^R until a hole \bar{s}_i , for $i \in I$, is reached. Thus if no hole is ever reached, the plays in the two games are identical, hence Player 2 wins following σ_2^T , as desired. If a hole $\bar{s}_i \in I$ is reached, then it is necessarily the case that $T_i \notin \text{lfp}(\mathbb{W})$, because otherwise the play in G_R would end in the state $\top = T_i \in \text{lfp}(\mathbb{W})$, which is winning for Player 1, while playing in accordance with the winning strategy σ_2^R . A contradiction. We then define the strategy σ_2^T to play the rest of the game (forgetting the previous history) as the strategy $\sigma_2^{T_i}$, constructed as for σ_2^T , but with respect to the branching play T_i . We shall call this a *re-starting* point of the play. We now prove that the strategy σ_2^T is winning for Player 2 as desired. There are two cases to consider. A play in the game \mathcal{G}_T played in accordance with σ_2^T can have,

- (1) either finitely many restarting points, or
- (2) infinitely many restarting points, i.e., infinitely many occurrences of states of the form $\langle p_i, X \rangle$, for $i \in I$.

In the first case, following earlier observations, the resulting path has a winning tail and thus is winning. In the second case, since the states of the form $\langle p_i, X \rangle$ are assigned maximal odd priority in $\mathcal{G}_\rho^{\mu X.F}$, the play is winning for Player 2 (see Definition 4.16). \square


The following results will be used for dealing with the measure-theoretic complications associated with the complexity of the winning set $\Phi^{\mu X.F}$.

Lemma 6.8. *The least fixed point of \mathbb{W} is reached in at most ω_1 iterations, i.e., the equality $\text{lfp}(\mathbb{W}) = \bigcup_{\alpha < \omega_1} \mathbb{W}_\alpha$ holds, where $\mathbb{W}_\alpha = \bigcup_{\beta < \alpha} \mathbb{W}(\mathbb{W}_\beta)$.*

Proof. Assume $T[T_i]_{i \in I} \in \mathbb{W}_{\omega_1+1}$ or, equivalently, $T[T_i \in \mathbb{W}_{\omega_1}] \in \Phi^F$. Let $J = \{i \mid T_i \in \mathbb{W}_{\omega_1}\}$. For each $j \in J$, let β_j be the least ordinal such that $T_j \in \mathbb{W}_{\beta_j}$. Note that β_j is a countable ordinal, since $\mathbb{W}_{\omega_1} = \bigcup_{\alpha < \omega_1} \mathbb{W}_\alpha$, and ω_1 is the least uncountable ordinal. Let $\beta = \bigsqcup_{j \in J} \beta_j$ be the supremum ordinal of all β_j . Note that β is a countable ordinal since I is countable and $J \subseteq I$. It then follows, by Definition 6.4, that $T[T_i \in \mathbb{W}_\beta]$ and $T[T_i \in \mathbb{W}_{\omega_1}]$ are identical branching plays. Thus $T[T_i]_{i \in I} \in \mathbb{W}_{\beta+1} \subseteq \mathbb{W}_{\omega_1}$. Hence $\mathbb{W}_{\omega_1+1} \subseteq \mathbb{W}_{\omega_1}$ as desired. \square

The result of Lemma 6.8 can be shown to be strict. One can indeed construct a $\text{pL}\mu^\odot$ game \mathcal{G}_ρ^F such that $\mathbb{W}_\alpha \subsetneq \Phi^{\mu X.F}$ for every countable ordinal α . We refer to Lemma 6.3.8 in [25] for a proof of this fact.

We now show that each set in the chain $\{\mathbb{W}_\alpha\}_{\alpha < \omega_1}$, having $\Phi^{\mu X.F}$ as limit, is provably universally measurable in $\text{ZFC} + \text{MA}_{\aleph_1}$ set theory.

Figure 4: Markov branching play $M[\lambda_i]_{i \in I}$

Lemma 6.9 (MA_{\aleph_1}). *If $X \subseteq \mathcal{BP}^{\mu X.F}$ is universally measurable then so is $\mathbb{W}(X)$, i.e., the function \mathbb{W} maps universally measurable sets to universally measurable sets.*

Proof. By Definition 6.4, we need to show that $\mathbb{W}(X) \stackrel{\text{def}}{=} m_X^{-1}(\Phi^F)$ is universally measurable. By application of Theorem 4.20 and Theorem 2.10, the winning set Φ^F is provably universally measurable in $\text{ZFC} + \text{MA}_{\aleph_1}$. The desired result then follows by application of Lemma 6.3 and Theorem 2.8(b). \square

The techniques adopted in this section can be trivially adapted to get the expected dual results which we simply summarize in the following proposition.

Proposition 6.10. *The following assertion holds for every $\text{pL}\mu^\odot$ formula $\nu X.F$ and interpretation ρ :*

- (1) *the winning set $\Phi^{\nu X.F}$ is the greatest fixed point of the operator \mathbb{W} ,*
- (2) *$\text{gfp}(\mathbb{W}) = \bigcap_{\alpha < \omega_1} \mathbb{W}^\alpha$, where $\mathbb{W}^\alpha = \bigcap_{\beta < \alpha} \mathbb{W}(\mathbb{W}^\beta)$,*
- (3) *(MA_{\aleph_1}) for every countable ordinal α , the set \mathbb{W}^α is universally measurable.*

7. ROBUST MARKOV BRANCHING PLAYS

In this section we identify a property of Markov branching plays in $\text{pL}\mu^\odot$ games which will be used in the proof of Theorem 5.5 in the next section.

Given a $\text{pL}\mu^\odot$ game G_ρ^F , with F a $\text{pL}\mu^\odot$ formula (having a free variable X) and ρ a $[0, 1]$ -interpretation of the variables, respectively, we extend the graphical notation, introduced in Section 6 for branching plays, to Markov branching plays in the expected way. Thus, given a Markov branching play M in G_ρ^F rooted at $\langle p, F \rangle$ (for some state p), we depict M as in Figure 4 and denote it by $M[\lambda_i]_{i \in I}$, with $\lambda_i \in [0, 1]$ being the probability labeling the edge connecting $\langle p_i, X \rangle$ with \top . Note that, by definition of the game G_ρ^F , the value λ_i coincides with $\rho(X)(p_i)$, for $i \in I$. However, it is convenient to consider, as a technical tool, Markov branching plays of the form $M[\lambda_i]_{i \in I}$ having λ_i , for $i \in I$, of an arbitrary value, even though these plays never really correspond to achievable plays in G_ρ^F . The associated probability measure $\mathbb{P}_{M[\lambda_i]_{i \in I}}$ over branching plays in G_ρ^F is defined as expected.

Definition 7.1 (Robust Markov branching plays). Fix a $\text{pL}\mu^\odot$ formula F (with a free variable X) and a $[0, 1]$ -interpretation ρ . We say that a Markov branching play $M[\lambda_i]_{i \in I}$ in G_ρ^F , for $I \subseteq \mathbb{N}$, is *robust in the variable X* , or just *robust* when X is clear from the context, if for every $\varepsilon > 0$ the following properties hold,

- (1) $\mathbf{E}(M[\gamma_i]) \geq \mathbf{E}(M[\lambda_i]) - \sum_{i \in I} \frac{\varepsilon}{2^{i+1}}$, and
- (2) $\mathbf{E}(M[\delta_i]) \leq \mathbf{E}(M[\lambda_i]) + \sum_{i \in I} \frac{\varepsilon}{2^{i+1}}$,




Figure 5: Illustrative example. The symbol $\overline{\lambda_i}$ denotes the value $1 - \lambda_i$.

for every sequences $\{\gamma_i\}_{i \in I}$ and $\{\delta_i\}_{i \in I}$ of reals in $[0, 1]$ such that, for every $i \in I$, the inequalities $\gamma_i \geq \lambda_i - \frac{\varepsilon}{\#(i)}$ and $\delta_i \leq \lambda_i + \frac{\varepsilon}{\#(i)}$ hold (see Definition 2.4 of $\# : \mathbb{N} \rightarrow \mathbb{N}$).

The notion of robustness captures a useful substitutivity property. If, in a Markov branching play $M[\lambda_i]_{i \in I}$, the probability λ_i of reaching from the state $\langle p_i, X \rangle$ the winning (for Player 1) state \top is replaced with a smaller but close enough value γ_i , then the resulting Markov branching play $M[\gamma_i]_{i \in I}$ has an expected value close to that of $M[\lambda_i]_{i \in I}$ too.


Note how, in Definition 7.1, the constraint on the distance between γ_i (δ_i) and λ_i crucially depends on the index $i \in I$. Definition 7.1 has been identified to meet the nature of Markov branching plays in $\text{pL}\mu^\odot$ games and, as we shall see in the next section, every Markov branching play in a $\text{pL}\mu^\odot$ game is indeed robust in every free variable. However, it is useful to observe that the dependence on $i \in I$ for the constraint on the distance between γ_i (δ_i) and λ_i is necessary. Indeed, the simpler property

$$(2') \quad \mathbf{E}(M[\delta_i]) \leq \mathbf{E}(M[\lambda_i]) + c_1 \quad \text{where } \forall i. (\delta_i < \lambda_i + c_2) \quad (7.1)$$

for some constant values $c_1, c_2 \in (0, 1)$, is not satisfied (in general) by $\text{pL}\mu^\odot$ Markov branching plays. In other words, Markov branching plays in $\text{pL}\mu^\odot$ are not stable (in their expected value) if the values λ_i are altered uniformly, i.e., by a fixed $c_2 > 0$. Consider, for example, the Markov branching play, having countably many holes, depicted as in Figure 5. Assume that $\lambda_i = 0$, for every $i \in \mathbb{N}$ and fix some $c_2 > 0$. Then it is simple to verify that $\mathbf{E}(M[\lambda_i]) = 0$ and $\mathbf{E}(M[\delta_i]) = 1$ (where $\delta_i = c_2$, for all $i \in I$), contradicting (2') above. This phenomenon reflects the discontinuity of the denotational interpretation of $\text{pL}\mu^\odot$ formulas on the free variables (see Proposition 3.5). Indeed note that the play of Figure 5 is a Markov branching play of a $\text{pL}\mu^\odot$ game associated with a formula of the form $\mu X.(Y \odot X) \stackrel{\text{def}}{=} \mathbb{P}_{>0} Y$, similar to the one discussed in the proof of Proposition 3.5.

We now establish a useful property relating expected values of Markov branching plays in \mathcal{G}_ρ^G and in the game $\mathcal{G}_\rho^{\mu X.G}$. Again, following established notation, we denote with $M[M_i]_{i \in I}$ the Markov branching play depicted as in Figure 6. The *Markov branching pre-play* $M[x_i]_{i \in I}$ captures the common structure of the Markov branching plays $M[M_i]_{i \in I}$ and $M[\lambda_i]_{i \in I}$ in $\mathcal{G}_\rho^{\mu X.F}$ and \mathcal{G}_ρ^F , respectively.

Lemma 7.2. *Let $\mu X.F$ be a $\text{pL}\mu^\odot$ formula and ρ a $[0, 1]$ -interpretation of the free variables. Let $M[M_i]_{i \in I}$ be a Markov branching play in $\mathcal{G}_\rho^{\mu X.F}$. For an ordinal β , let $\gamma_i^\beta =$*

Figure 6: Markov Branching play $M[M_i]_{i \in I}$

$\mathbb{P}_{M_i}(\bigcup_{\alpha < \beta} \mathbb{W}_\alpha)$ be the probability of the event $\bigcup_{\alpha < \beta} \mathbb{W}_\alpha$ (see Lemma 6.8) associated with the sub-Markov branching play M_i , for $i \in I$. Then the equality

$$\mathbb{P}_{M[M_i]_{i \in I}}(\mathbb{W}_\beta) = \mathbb{P}_{M[\gamma_i^\beta]_{i \in I}}(\Phi^F)$$

holds, where Φ^F is the winning set of the game \mathcal{G}_ρ^F .

Proof. Consider the function m_X , as specified in Definition 6.2. Recall that, by definition, the equalities

$$\mathbb{W}_\beta = \bigcup_{\alpha < \beta} \mathbb{W}(\mathbb{W}_\alpha) = m_{\bigcup_{\alpha < \beta} \mathbb{W}_\alpha}^{-1}(\Phi^F)$$

hold. The proof is completed by showing that the following property holds:

$$\mathbb{P}_{M[M_i]_{i \in I}}(m_{\bigcup_{\alpha < \beta} \mathbb{W}_\alpha}^{-1}(X)) = \mathbb{P}_{M[\gamma_i]_{i \in I}}(X), \quad (7.2)$$

for every Borel measurable set $X \subseteq \mathcal{BP}^F$, where \mathcal{BP}^F denotes the set of branching plays in \mathcal{G}_ρ^F . Indeed the desired result follows from Equation 7.2 by taking $X = \Phi^F$. Since probability measures in Polish spaces are regular, we can restrict X to range over basic open sets. Equation 7.2 can then be proved, with routine techniques, by induction of the complexity of basic open sets O_T (see Definition 4.5), i.e., on the size of the finite tree T . \square

Note that, as observed earlier, the Markov branching play $M[\gamma_i]_{i \in I}$ considered in Lemma 7.2, might not be a real play in \mathcal{G}_ρ^F , i.e., one induced by a strategy profile. This is the case when $\gamma_i \neq \rho(X)(p_i)$ for some $i \in I$.

8. PROOF OF EQUIVALENCE OF GAME AND DENOTATIONAL SEMANTICS

This section is devoted to the proof of Theorem 5.5. The proof technique we adopt is based on the *unfolding method* of [11]. The unfolding method can be roughly described as a technique for proving *properties* of (some sort of) two-player parity games by induction on the number of priorities used in the game. Usually, the first step is to prove that the property under consideration holds for all parity games with just one priority. Then the general result for games with $n + 1$ priorities follows by some argument making use of the inductive hypothesis. In our setting we are interested in two-player meta-parity games of the form \mathcal{G}_ρ^F , and the property we want to prove is that the lower and upper values of these games coincide with the denotational value of F under the interpretation ρ . We prove this by induction of the structure of F rather than on the number of priorities used in the game \mathcal{G}_ρ^F . This allows a more transparent and arguably simpler proof. The structure of our proof closely resembles the one of [26], where the equivalence of the game and denotational

semantics for the logic $\text{pL}\mu$ is proved using the unfolding method. In the present context, the proof is a significantly more technical undertaking due to the complexity of two-player meta-parity games: the Δ_2^1 -complexity of the winning sets (Theorem 4.20), their transfinite inductive characterization up to the first uncountable ordinal ω_1 (Theorem 6.8) and the discontinuity in the free variables of the denotational semantics (see Theorem 3.5 and Section 7). In what follows, we shall focus primarily on the novel aspects of the proof, referring to [26] for a detailed analysis of the simpler cases that easily generalize to the present setting.

We prove, by induction on the structure of the formulas that, for every PLTS $\mathcal{L} = \langle P, \{\xrightarrow{a}\}_{a \in L} \rangle$, $\text{pL}\mu^\odot$ formula F and $[0, 1]$ -interpretation ρ of the variables, the following assertions hold:

$$\llbracket F \rrbracket_\rho(p) = \text{VAL}_\downarrow(\mathcal{G}_\rho^F)(\langle p, F \rangle) = \text{VAL}_\uparrow(\mathcal{G}_\rho^F)(\langle p, F \rangle), \quad (8.1)$$

and

$$M \text{ is robust in the variable } X \text{ (see Definition 7.1)} \quad (8.2)$$

for every free variable X in F , and Markov branching play M rooted at $\langle p, F \rangle$ in \mathcal{G}_ρ^F .

Base case: $F = X$, for some variable $X \in \text{Var}$.

It follows immediately by application of Theorem 4.22 that point 8.1 holds. For what concerns point 8.2, a Markov branching play M rooted at $\langle p, X \rangle$ in \mathcal{G}_ρ^G is of the following form:

$$\begin{array}{c} \top \\ \lambda_0 \searrow \quad \nearrow 1 - \lambda_0 \\ \langle p_0, X \rangle \end{array}$$

where $\lambda_0 = \rho(X)(p)$. Thus, M has only one hole (i.e., $M = M[\lambda_0]$) and $\mathbf{E}(M) = \lambda_0$. The Markov branching play $M[\gamma]$, can then be depicted as above, replacing λ_0 with γ , for every $\gamma \in [0, 1]$. Therefore, $\mathbf{E}(M[\gamma]) = \gamma$ and point 8.2 is trivially satisfied as desired.

Inductive cases $F = G_1 \star G_2$ or $F = \circ G$, for $\star \in \{\vee, \wedge, \cdot, \odot\}$ and $\circ \in \{\langle a \rangle, [a]\}$.

For all these cases, the proof of point 8.1 follows easily by application of Theorem 4.22. The result can be proved following the same lines of the proof of (the corresponding inductive cases of) Theorem 5.1 in [26], thus we omit the routine details.

We now prove that point 8.2 holds for $F = G_1 \cdot G_2$. The other cases can be proved in a similar way. Let us consider an arbitrary Markov branching play M in \mathcal{G}_ρ^F rooted at $\langle p, G_1 \cdot G_2 \rangle$. Then, M can be depicted as follows, where M_1 and M_2 denote the two sub-Markov branching plays rooted at $\langle p, G_1 \rangle$ and $\langle p, G_2 \rangle$.

$$\begin{array}{ccc} & M_1 & M_2 \\ & \text{---} & \text{---} \\ & \swarrow \quad \searrow & \swarrow \quad \searrow \\ \langle p, G_1 \rangle & & \langle p, G_2 \rangle \\ & \swarrow \quad \searrow & \\ & \langle p, G_1 \cdot G_2 \rangle & \end{array}$$

Note that, by definition of the game \mathcal{G}_ρ^F , the sub-Markov branching play M_i is also a Markov branching play rooted at $\langle p, G_i \rangle$ in the game $\mathcal{G}_\rho^{G_i}$, for $i \in \{1, 2\}$. Also note that X is free in both G_1 and G_2 , since it is free in G by assumption. Let $M[x_i]_{i \in I}$ be the Markov branching pre-play obtained by pruning M at the states of the form $\langle p_i, X \rangle$, as described in Section

7. Let λ_i , for $i \in I$, be the probability labeling the edge connecting the i -th hole in M (i.e., the state $\langle p_i, X \rangle$), to the state \top . By Definition 7.1, we need to prove that

- (1) $\mathbf{E}(M[\gamma_i]_{i \in I}) \geq \mathbf{E}(M[\lambda_i]_{i \in I}) - \sum_{i \in I} \frac{\epsilon}{2^{i+1}}$, and
- (2) $\mathbf{E}(M[\delta_i]_{i \in I}) \leq \mathbf{E}(M[\lambda_i]_{i \in I}) + \sum_{i \in I} \frac{\epsilon}{2^{i+1}}$,

hold, for every sequence $\{\gamma_i\}_{i \in I}$ and $\{\delta_i\}_{i \in I}$ of reals in $[0, 1]$ such that, for every $i \in I \subseteq \mathbb{N}$, the inequalities $\gamma_i \geq \lambda_i - \frac{\epsilon}{\#(i)}$ and $\delta_i \leq \lambda_i + \frac{\epsilon}{\#(i)}$ hold. We just show how to prove the first inequality because the second one can be proved in a similar way.

Let I_1 and I_2 be the partition of the index set I specified as follows. The index $i \in I$ of a path \vec{s}_i in $M[x_i]_{i \in I}$ (connecting $\langle p, G_1 \odot G_2 \rangle$ with the hole $\langle p_i, X \rangle$) is in I_1 if \vec{s}_i lies in M_1 , i.e., its second state is $\langle p, G_1 \rangle$. Similarly, the index i belongs to I_2 if \vec{s}_i lies in M_2 . Note that I_1 indexes all the holes of the sub-Markov branching play M_1 and, similarly, I_2 indexes the holes of M_2 . Thus, let us denote with $M_1[x_i]_{i \in I_1}$ the Markov branching pre-play associated with M_1 , and similarly for $M_2[x_j]_{j \in I_2}$ and M_2 . Then, by the inductive hypothesis on the complexity of G_1 and G_2 , we know that M_1 and M_2 are robust in X , i.e., the following assertions, with respect to inequality (1) above,

- a) $\mathbf{E}(M_1[\gamma_i]_{i \in I_1}) \geq \mathbf{E}(M_1[\lambda_i]_{i \in I_1}) - \sum_{i \in I_1} \frac{\epsilon}{2^{i+1}}$, and
- b) $\mathbf{E}(M_2[\gamma_j]_{j \in I_2}) \geq \mathbf{E}(M_2[\lambda_j]_{j \in I_2}) - \sum_{j \in I_2} \frac{\epsilon}{2^{j+1}}$,

hold. By applying the product measure technique adopted in the proof of Theorem 4.22, it is easy to verify that the equalities $\mathbf{E}(M[\lambda_i]_{i \in I}) = \mathbf{E}(M_1[\lambda_i]_{i \in I_1}) \cdot \mathbf{E}(M_2[\lambda_j]_{j \in I_2})$ and $\mathbf{E}(M[\gamma_i]_{i \in I}) = \mathbf{E}(M_1[\gamma_i]_{i \in I_1}) \cdot \mathbf{E}(M_2[\gamma_j]_{j \in I_2})$ hold. The desired equation (1) above can then be derived as follows:

$$\begin{aligned}
\mathbf{E}(M[\gamma_i]_{i \in I}) &= \mathbf{E}(M_1[\gamma_i]_{i \in I_1}) \cdot \mathbf{E}(M_2[\gamma_j]_{j \in I_2}) \\
&\geq_b \mathbf{E}(M_1[\gamma_i]_{i \in I_1}) \cdot (\mathbf{E}(M_2[\lambda_j]_{j \in I_2}) - \sum_{j \in I_2} \frac{\epsilon}{2^{j+1}}) \\
&\geq_* (\mathbf{E}(M_1[\gamma_i]_{i \in I_1}) \cdot \mathbf{E}(M_2[\lambda_j]_{j \in I_2})) - \sum_{j \in I_2} \frac{\epsilon}{2^{j+1}} \\
&\geq_a \left((\mathbf{E}(M_1[\lambda_i]_{i \in I_1}) - \sum_{i \in I_1} \frac{\epsilon}{2^{i+1}}) \cdot \mathbf{E}(M_2[\lambda_j]_{j \in I_2}) \right) - \sum_{j \in I_2} \frac{\epsilon}{2^{j+1}} \\
&\geq_* (\mathbf{E}(M_1[\lambda_i]_{i \in I_1}) \cdot \mathbf{E}(M_2[\lambda_j]_{j \in I_2})) - \sum_{i \in I_1} \frac{\epsilon}{2^{i+1}} - \sum_{j \in I_2} \frac{\epsilon}{2^{j+1}} \\
&\geq_{I=I_1 \uplus I_2} (\mathbf{E}(M_1[\lambda_i]_{i \in I_1}) \cdot \mathbf{E}(M_2[\lambda_j]_{j \in I_2})) - \sum_{i \in I} \frac{\epsilon}{2^{i+1}} \\
&= \mathbf{E}(M[\lambda_i]_{i \in I}) - \sum_{i \in I} \frac{\epsilon}{2^{i+1}}
\end{aligned}$$

where the steps labeled with $(*)$ are valid because all terms have values in the interval $[0, 1]$.

Inductive case $\mathbf{F} = \mu \mathbf{X}. \mathbf{G}$.

- We first prove that point 8.1 holds.

For every state p and every interpretation ρ we have, by definition of the denotational semantics, that the following equality holds:

$$\llbracket \mu X. G \rrbracket_\rho(p) \stackrel{\text{def}}{=} \text{lfp} \left(\lambda f \in [0, 1]^P. (\llbracket G \rrbracket_{\rho[f/X]}) \right)(p).$$

By the Knaster–Tarski theorem, the previous equation can be rewritten as:

$$\llbracket \mu X. G \rrbracket_\rho(p) = \bigsqcup_{\alpha} \llbracket G \rrbracket_\rho^\alpha, \quad (8.3)$$

where α ranges over the ordinals, and $\llbracket G \rrbracket_\rho^\alpha$ is defined as $\bigsqcup_{\beta < \alpha} \llbracket G \rrbracket_{\rho[\llbracket G \rrbracket_\rho^\beta/X]}$. Let us denote with γ the least ordinal such that $\llbracket G \rrbracket_\rho^\gamma = \llbracket \mu X. G \rrbracket_\rho$, and with $\rho_\gamma \in [0, 1]^P$ the interpretation $\rho[\llbracket G \rrbracket_\rho^\gamma/X]$. Thus, the following equation holds:

$$\llbracket G \rrbracket_{\rho_\gamma} = \llbracket \mu X. G \rrbracket_\rho. \quad (8.4)$$

Let us now turn our attention to the $2\frac{1}{2}$ -player meta-parity game $\mathcal{G}_\rho^{\mu X.G}$. Our goal is to prove that point 8.1 holds, i.e., that the following equalities

$$\llbracket \mu X.G \rrbracket_\rho(p) = \text{VAL}_\downarrow(\mathcal{G}_\rho^{\mu X.G})(\langle p, \mu X.G \rangle) = \text{VAL}_\uparrow(\mathcal{G}_\rho^{\mu X.G})(\langle p, \mu X.G \rangle) \quad (8.5)$$

hold, for every $p \in P$. As a first observation, note that the state $\langle p, \mu X.G \rangle$ is not reachable by any other game state, and that it has the state $\langle p, G \rangle$ as its only successor state. It then follows by application of Proposition 4.22 that, in order to prove the desired result (8.5), we just have to show that the equalities

$$\llbracket G \rrbracket_{\rho_\gamma}(p) = \text{VAL}_\downarrow(\mathcal{G}_\rho^{\mu X.G})(\langle p, G \rangle) = \text{VAL}_\uparrow(\mathcal{G}_\rho^{\mu X.G})(\langle p, G \rangle) \quad (8.6)$$

hold. In order to improve readability, we shall denote with $\llbracket \mu X.G \rrbracket_\rho^*: P \rightarrow [0, 1]$ the function defined as $\lambda p \in P. (\text{VAL}_\star(\mathcal{G}_\rho^{\mu X.G})(\langle p, G \rangle))$, for $\star \in \{\downarrow, \uparrow\}$. Thus, Equation 8.6 can be rewritten as follows:

$$\llbracket G \rrbracket_{\rho_\gamma}(p) = \llbracket \mu X.G \rrbracket_\rho^\downarrow(p) = \llbracket \mu X.G \rrbracket_\rho^\uparrow(p). \quad (8.7)$$

Note that the analogous functions $\llbracket G \rrbracket_{\rho[f/X]}^\star: P \rightarrow [0, 1]$ specified, for $\star \in \{\downarrow, \uparrow\}$, as $\lambda p \in P. (\text{Val}_\star(\mathcal{G}_{\rho[f/X]}^G)(\langle p, G \rangle))$, satisfy the following equation:

$$\llbracket G \rrbracket_{\rho[f/X]} = \llbracket G \rrbracket_{\rho[f/X]}^\downarrow = \llbracket G \rrbracket_{\rho[f/X]}^\uparrow \quad (8.8)$$

for all $f \in [0, 1]^P$, by induction hypothesis (8.1) on G .

We prove Equation 8.7 by exploiting the similarities between the game $\mathcal{G}_\rho^{\mu X.G}$ and the game $\mathcal{G}_{\rho[f/X]}^G$, already discussed in Section 6. Our first observation is the following:

$$\llbracket \mu X.G \rrbracket_\rho^* = \llbracket G \rrbracket_{\rho[\llbracket \mu X.G \rrbracket_\rho^*/X]}^* \stackrel{\text{Eq. 8.8}}{=} \llbracket G \rrbracket_{\rho[\llbracket \mu X.G \rrbracket_\rho^*/X]} \quad (8.9)$$

for $\star \in \{\downarrow, \uparrow\}$. Indeed when a state of the form $\langle p, X \rangle$ is reached in the game $\mathcal{G}_\rho^{\mu X.G}$ the play continues from the state $\langle p, G \rangle$ and ends in a victory for Player 1 with (lower or upper) value $\llbracket \mu X.G \rrbracket_\rho^*(p)$, and, similarly, when the state $\langle p, X \rangle$ is reached in $\mathcal{G}_{\rho[\llbracket \mu X.G \rrbracket_\rho^*/X]}^G$, the play immediately terminates in favor of Player 1 with probability $\llbracket \mu X.G \rrbracket_\rho^*(p)$.

By application of Equation 8.8, this implies that both $\llbracket \mu X.G \rrbracket_\rho^\uparrow$ and $\llbracket \mu X.G \rrbracket_\rho^\downarrow$ are fixed points of the functional $\lambda f \in [0, 1]^P. (\llbracket G \rrbracket_{\rho[f/X]})$. Note that, for all $p \in P$, the inequality $\llbracket \mu X.G \rrbracket_\rho^\downarrow(p) \leq \llbracket \mu X.G \rrbracket_\rho^\uparrow(p)$ trivially holds. Moreover the inequality $\llbracket \mu X.G \rrbracket_\rho(p) \leq \llbracket \mu X.G \rrbracket_\rho^\uparrow(p)$ holds, for all $p \in P$, because $\llbracket \mu X.G \rrbracket_\rho$ (or, equivalently, $\llbracket G \rrbracket_{\rho_\gamma}$) is the least fixed point of $\lambda f \in [0, 1]^P. (\llbracket G \rrbracket_{\rho[f/X]})$.

We shall prove the desired result (Equation 8.7) by showing that, for all $p \in P$, the inequality

$$\llbracket \mu X.G \rrbracket_\rho^\uparrow(p) \stackrel{\text{def}}{=} \text{Val}_\uparrow(\mathcal{G}_\rho^{\mu X.G})(\langle p, G \rangle) \leq \llbracket G \rrbracket_{\rho_\gamma}(p) \quad (8.10)$$

holds. We do this by constructing, for every $\varepsilon > 0$ a strategy σ_2^ε for Player 2 in the game $\mathcal{G}_\rho^{\mu X.G}$, satisfying the following inequality:

$$\bigsqcup_{\sigma_1} \mathbf{E}(M_{\sigma_1, \sigma_2^\varepsilon}^{\langle p, G \rangle}) \leq \llbracket G \rrbracket_{\rho_\gamma}(p) + \varepsilon. \quad (8.11)$$

The strategy σ_2^ε is constructed using the collection of δ -optimal strategies, for $\delta > 0$, in the game $\mathcal{G}_{\rho_\gamma}^G$, i.e., strategies τ_2^δ such that the following equality holds:

$$\bigsqcup_{\tau_1} \mathbf{E}(M_{\tau_1, \tau_2^\delta}^{\langle p, G \rangle}) \leq \text{VAL}(\mathcal{G}_{\rho_\gamma}^G)(\langle p, G \rangle) + \delta \stackrel{\text{Eq. 8.8}}{=} \llbracket G \rrbracket_{\rho_\gamma}(p) + \delta. \quad (8.12)$$

Let $e: \mathcal{P}_{\mu X.G}^{<\omega} \rightarrow \mathbb{N}$ be a numbering (i.e., an injective map into the natural numbers) of the finite paths in the game $\mathcal{G}_\rho^{\mu X.G}$. The strategy σ_2^ε is defined, for every $\varepsilon > 0$, as follows:

$$\sigma_2^\varepsilon(\vec{s}) = \begin{cases} \tau_2^{\frac{\varepsilon}{2}}(\vec{s}) & \text{if } \vec{s} \text{ does not contain states of the form } \langle p, X \rangle, \text{ for } p \in P \\ \sigma_2^{\frac{\varepsilon}{2} \cdot \frac{1}{\#(e(\vec{s}_j))}}(\vec{t}) & \text{if } \vec{s} = \vec{s}_j.\vec{t} \text{ with } \text{last}(\vec{s}_j) = \langle p, X \rangle \text{ for some } p \in P \end{cases}$$

for every finite path \vec{s} whose last state belong to Player 2 (i.e., such that $\text{last}(\vec{s}) \in S_2$), where the function $\#: \mathbb{N} \rightarrow \mathbb{N}$ is specified as in Definition 2.4. The strategy σ_2^ε can be informally described as follows. When the game begins σ_2^ε behaves as the strategy $\tau_2^{\frac{\varepsilon}{2}}$ until a state of the form $\langle p, X \rangle$ is reached. This is a good definition because plays in the two games $\mathcal{G}_\rho^{\mu X.G}$ and $\mathcal{G}_{\rho_\gamma}^G$ are identical until states of this form are reached. If a state of the form $\langle p, X \rangle$ is eventually reached following some path \vec{s}_j , then Player 2 *improves* their strategy, and plays the rest of the game (starting at the unique successor state $\langle p, G \rangle$ of $\langle p, X \rangle$) in accordance with the better strategy $\sigma_2^{\frac{\varepsilon}{2} \cdot \frac{1}{\#(e(\vec{s}_j))}}$. Note how the choice of the new strategy to be followed crucially depends on the path \vec{s}_j .

We are now going to show that, for every $\varepsilon > 0$, the strategy σ_2^ε satisfies the desired inequality 8.11. We need to show that, for every strategy σ_1 for Player 1, the inequality

$$\mathbf{E}(M_{\sigma_1, \sigma_2^\varepsilon}^{\langle p, G \rangle}) \leq \llbracket G \rrbracket_{\rho_\gamma}(p) + \varepsilon. \quad (8.13)$$

holds. Recall that, by definition, $\mathbf{E}(M_{\sigma_1, \sigma_2^\varepsilon}^{\langle p, G \rangle}) = \mathbb{P}_{\sigma_1, \sigma_2^\varepsilon}^s(\Phi)$, where $\mathbb{P}_{\sigma_1, \sigma_2^\varepsilon}^s$ denotes the probability measure over branching plays in $\mathcal{G}_\rho^{\mu X.G}$ induced by the Markov branching play $M_{\sigma_1, \sigma_2^\varepsilon}^s$, and Φ denotes the set of winning branching plays for Player 1 in $\mathcal{G}_\rho^{\mu X.G}$. By Theorem 6.7 and Lemma 6.8, we know that $\Phi = \bigcup_{\alpha < \omega_1} \mathbb{W}_\alpha$. Thus, the previous equation can be rewritten as follows:

$$\mathbb{P}_{\sigma_1, \sigma_2^\varepsilon}^{\langle p, G \rangle} \left(\bigcup_{\alpha < \omega_1} \mathbb{W}_\alpha \right) \leq \llbracket G \rrbracket_{\rho_\gamma}(p) + \varepsilon. \quad (8.14)$$



By application of Theorem 2.10, under the set-theoretic assumption MA_{\aleph_1} we can further rewrite the previous equation as follows:

$$\bigsqcup_{\alpha < \omega_1} (\mathbb{P}_{\sigma_1, \sigma_2^\varepsilon}^{\langle p, G \rangle}(\mathbb{W}_\alpha)) \leq \llbracket G \rrbracket_{\rho_\gamma}(p) + \varepsilon. \quad (8.15)$$

This step allow us to set up a proof by transfinite induction for the desired Equality 8.11. We shall now prove that for every countable ordinal $\beta < \omega_1$ and every $\varepsilon > 0$, the inequality

$$\mathbb{P}_{\sigma_1, \sigma_2^\varepsilon}^{\langle p, G \rangle}(\mathbb{W}_\beta) \leq \llbracket G \rrbracket_{\rho_\gamma}(p) + \varepsilon. \quad (8.16)$$

holds. Assume Equation 8.16 holds for every ordinal $\alpha < \beta$. The Markov branching play $M_{\sigma_1, \sigma_2^\varepsilon}^{\langle p, G \rangle}$ can be depicted, following the notation introduced in sections 6 and 7, as in Figure 7. By definition of σ_2^ε , the sub-Markov branching plays M_i , for $i \in I$, are Markov branching plays played in accordance with the strategy $\sigma_2^{\varepsilon_i}$, with $\varepsilon_i = \frac{\varepsilon}{2} \cdot \frac{1}{\#(e(\vec{s}_i))}$. Let us denote with

Figure 7: Markov Branching play $M[M_i]_{i \in I}$ Figure 8: Markov branching play $M[\delta_i]_{i \in I}$

δ_i^α the value $\mathbb{P}_{M_i}(\mathbb{W}_\alpha)$, where \mathbb{P}_{M_i} denotes the probability measure associated with M_i , and $\alpha < \beta$. Moreover, let us denote with δ_i , for $i \in I$ the value $\bigsqcup_{\alpha < \beta} \delta_i^\alpha$. Then, by induction hypothesis on the ordinals, we know that, for every $\alpha < \beta$, the inequality

$$\delta_i^\alpha \leq \llbracket G \rrbracket_{\rho_\gamma}(p) + \frac{\varepsilon}{2} \cdot \frac{1}{\#(e(\vec{s}_i))} \quad (8.17)$$

holds, for every $i \in I$. By application of Lemma 7.2, we know that the following equality holds:

$$\mathbb{P}_{\sigma_1, \sigma_2^\varepsilon}^{\langle p, G \rangle}(\mathbb{W}_\beta) = \mathbf{E}(M[\delta_i]_{i \in I}) \quad (8.18)$$

where $M[\delta_i]_{i \in I}$ denotes the Markov branching play in the game $\mathcal{G}_{\rho_\gamma}^G$, obtained from $M_{\sigma_1, \sigma_2^\varepsilon}^{\langle p, G \rangle}$ as specified in Section 7, which can be depicted as in Figure 8. Let us denote with λ_i , for $i \in I$, the value $\llbracket G \rrbracket_{\rho_\gamma}(p_i)$. Observe that the Markov branching play $M[\lambda_i]_{i \in I}$ (depicted by replacing δ_i with λ_i in Figure 8) is a Markov branching play in the game $\mathcal{G}_{\rho_\gamma}^G$. Moreover, by definition of σ_2^ε , the Markov branching play $M[\lambda_i]_{i \in I}$ is played by Player 2 in accordance with the $\frac{\varepsilon}{2}$ -optimal strategy $\tau_2^{\frac{\varepsilon}{2}}$. It then follows by Equation 8.12 that

$$\mathbf{E}(M[\lambda_i]_{i \in I}) \leq \llbracket G \rrbracket_{\rho_\gamma}(p) + \frac{\varepsilon}{2} \quad (8.19)$$

Recall that, by induction hypothesis (8.2) on G , the Markov branching play $M[\lambda_i]_{i \in I}$ is robust is the free (in G) variable X . Thus the following inequality holds:

$$\mathbf{E}(M[\delta_i]_{i \in I}) \leq \mathbf{E}(M[\lambda_i]_{i \in I}) + \sum_{i \in I} \left(\frac{\varepsilon}{2} \cdot \frac{1}{\#(e(\vec{s}_i))} \right). \quad (8.20)$$

Since the numbering e is injective, it follows from equation 8.19 and 8.20 that $\mathbf{E}(M[\delta_i]_{i \in I}) \leq \llbracket G \rrbracket_{\rho_\gamma}(p) - \frac{\varepsilon}{2} - \frac{\varepsilon}{2}$. Thus, Equation 8.16 holds as desired.

- We now prove that point 8.2 holds.

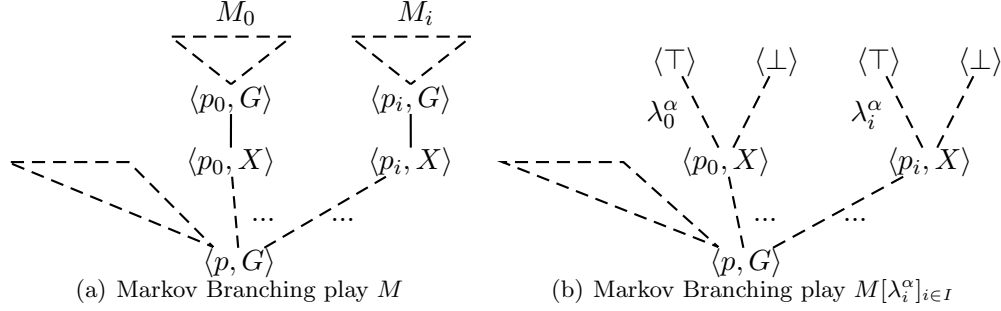


Figure 9: Markov branching plays $M = M[M_i]_{i \in I}$ in $\mathcal{G}_\rho^{\mu X.G}$ and $M[\lambda_i^\alpha]_{i \in I}$ in \mathcal{G}_ρ^G

We just discuss the main ideas of the proof as the necessary techniques have been already introduced in the proof of point 8.1 above.

We need to show that every Markov branching play M in $\mathcal{G}_\rho^{\mu X.G}$, rooted at $\langle p, \mu X.G \rangle$ is robust in every free variable Y . Clearly $Y \neq X$, since X is bound in $\mu X.F$. As already observed earlier, we can just consider Markov branching plays M rooted at $\langle p, G \rangle$, since the state $\langle p, \mu X.G \rangle$ has $\langle p, G \rangle$ as its only successor state, and it is not reachable by any other states. Let $J \subseteq \mathbb{N}$ be the set indexing the paths from the root $\langle p, G \rangle$ of M to states of the form $\langle p_j, Y \rangle$. Let us denote with $M[y_j]_{j \in J}$ the associated Markov branching pre-play. Let $\lambda_j \in [0, 1]$, for $j \in J$, be the value labeling the edge connecting the j -th hole in M with the state \top , i.e., $\lambda_j = \rho(Y)(p_j)$. Thus $M = M[\lambda_j]_{j \in J}$. We need to prove that, for every $\varepsilon > 0$, the inequalities

- (1) $\mathbf{E}(M[\gamma_j]_{j \in J}) \geq \mathbf{E}(M[\lambda_j]_{j \in J}) - \sum_{j \in J} \frac{\varepsilon}{2^{j+1}}$, and
- (2) $\mathbf{E}(M[\delta_j]_{j \in J}) \leq \mathbf{E}(M[\lambda_j]_{j \in J}) + \sum_{j \in J} \frac{\varepsilon}{2^{j+1}}$,

hold for every sequences $\{\gamma_j\}_{j \in J}$ and $\{\delta_j\}_{j \in J}$ of reals in $[0, 1]$ such that, for every $j \in J \subseteq \mathbb{N}$, the inequalities $\gamma_j \geq \lambda_j - \frac{\varepsilon}{\#(j)}$ and $\delta_j \leq \lambda_j + \frac{\varepsilon}{\#(j)}$ hold. We just consider the first inequality, as the second one can be proved in a similar way.

Recall that, by application of Theorem 2.10, under the set-theoretic assumption MA_{\aleph_1} , the equality $\mathbf{E}(M[\gamma_j]_{j \in J}) \stackrel{\text{def}}{=} \mathbb{P}_{M[\gamma_j]_{j \in J}}(\Phi) = \bigsqcup_{\alpha < \omega_1} \mathbb{P}_{M[\gamma_j]_{j \in J}}(\mathbb{W}_\alpha)$ holds. Similarly for $\mathbf{E}(M[\lambda_j]_{j \in J})$. We shall then prove (1) above, by proving the following more general property: for all Markov branching plays $M = M[\lambda_j]_{j \in J}$ in $\mathcal{G}_\rho^{\mu X.G}$ rooted at $\langle p, G \rangle$ and sequence $\{\gamma_j\}_{j \in J}$ as described above, the following equality holds:

$$\mathbb{P}_{M[\gamma_j]_{j \in J}}(\mathbb{W}_\alpha) \geq \mathbb{P}_{M[\lambda_j]_{j \in J}}(\mathbb{W}_\alpha) - \sum_{j \in J} \frac{\varepsilon}{2^{j+1}} \quad (8.21)$$

This is proven, again, by transfinite induction on the ordinals. Suppose the property holds for all $\alpha < \beta$. The main idea to prove the inductive step is to reduce the problem on M (which can be depicted as in Figure 9(a) by exposing the collection of paths reaching states of the form $\langle p_i, X \rangle$, for an index set I) to that of $M[\lambda_i^\alpha]_{i \in I}$ (where $\lambda_i^\alpha = \mathbb{P}_{M_i}(\bigcup_{\alpha < \beta} \mathbb{W}_\alpha)$), the Markov branching play in \mathcal{G}_ρ^G constructed as described in Lemma 7.2 and depicted in Figure 9(b). The result then follows by applications of Lemma 7.2 (relating the values $\mathbb{P}_M(\mathbb{W}_\beta)$ and of $\mathbf{E}(\mathbb{P}_{M[\lambda_i^\alpha]})$), by induction hypothesis on α (8.21) and by induction hypothesis (8.2) on the robustness of the Markov branching plays $M[\lambda_i^\alpha]_{i \in I}$ in \mathcal{G}_ρ^G . We omit the technical

details. A detailed proof of (a generalized version of) this result can be found in [25, §4.3].

Inductive case $G = \nu X.G$.

Similar to the previous one and based on the properties summarized in Proposition 6.10.

8.1. Remarks about the use of MA_{\aleph_1} . We conclude this section by highlighting the two critical steps in our proof where the set-theoretic assumption MA_{\aleph_1} is used.

As already discussed in Section 4, we make use of Martin's Axiom at \aleph_1 to ensure that the Δ_2^1 winning set of a $pL\mu^\odot$ game is universally measurable. The universal measurability of Δ_2^1 sets can, however, be proved in other extensions of ZFC. For example, determinacy-based axioms such as *Analytic Determinacy* [15], suffice (see, e.g., [17, §36.E]).

The second use we make of the axiom MA_{\aleph_1} is in the derivation of Equation 8.15 in the proof above. This is a fundamental step required to set up a proof by transfinite induction on the countable ordinals. As stated in Theorem 2.10, one of the consequences of Martin's Axiom at \aleph_1 is that probability measures on Polish spaces are ω_1 -continuous. Such a property, clearly implying the negation of the Continuum hypothesis, does not follow from determinacy-based axioms mentioned above. As mentioned after Lemma 6.8, our inductive characterization of the winning set of $pL\mu^\odot$ games, can be shown to require, in general, ω_1 -iterations (i.e., approximants) to converge. Thus some form of ω_1 -continuity seems to be required by the proof technique adopted in this Section.

9. CONCLUSIONS AND FUTURE WORK

One of the primary interests in a game semantics for $pL\mu^\odot$, and more generally for all logics having a $[0, 1]$ -valued semantics with an intended probabilistic reading, is to offer an accessible and clear interpretation for the property described by a formula. We suggest that our game semantics, built on top of the elementary idea of concurrent execution of independent sub-instances of the game, succeeds in this task. The logic $pL\mu^\odot$ is expressive enough to encode the qualitative threshold modalities $\mathbb{P}_{>0}$ and $\mathbb{P}_{=1}$ which, as discussed at the end of Section 5, are interpreted within the game semantics in a straightforward way.

Despite the naturalness of the definition, our proof of equivalence is a technically undertaking. Our result, based on a transfinite characterization of $pL\mu^\odot$ winning sets, is carried out in $ZFC + MA_{\aleph_1}$ set-theory to deal with the measure theoretic complications associated with the complexity of winning sets. We are not aware of any other result in theoretical computer science whose proof is (or at least was originally) carried out in proper extensions of ZFC set theory. Thus, our proof of determinacy, which is generalized in [25] from the $pL\mu^\odot$ games considered in this paper to arbitrary two-player stochastic meta-parity games, is perhaps noteworthy as being a first example of this kind of result.

Although the logic $pL\mu^\odot$ subsumes the qualitative fragment of the logic PCTL of [2], it does not seem possible to encode the full logic PCTL within $pL\mu^\odot$. In [25] an extension of $pL\mu^\odot$, obtained by adding to the syntax of the logic yet another pair of De Morgan dual connectives, is considered. Interestingly, this extension, which is capable of encoding the full logic PCTL, can be given an appropriate game semantics in terms of two-player stochastic meta-parity games. We refer to [25] for a proof of this fact, which serves as a demonstration of the expressive power of the new class of two-player stochastic meta-parity games introduced in this paper.

Our work leaves open several directions for future research. From a theoretical point of view, it would be interesting to remove the dependencies on the set-theoretic assumption MA_{\aleph_1} from our proof. This, in light of the remarks of Section 8.1, looks as a challenging task. In another direction, it would be interesting to investigate the theory of two-player stochastic meta-games. Preliminary results, such as the fact that Blackwell games can be encoded as tree games and that the open problem of *qualitative determinacy* [6] for stochastic games can be formulated as a determinacy problem for tree games, are obtained in the author's PhD thesis [25]. However, several questions remain open. For example, the concept of independent execution of actions in two-player games have been already considered as a tool for understanding logics of imperfect information (see, e.g., [4], [12]). It is then be natural to explore the expressive power of tree games in this setting. A very general class of two-player games with concurrent behaviors, based on event structures, has recently been considered in [9] where a determinacy theorem is obtained for games satisfying appropriate restrictions. Comparing our notion of tree games with their *concurrent games*, and the corresponding determinacy results, looks like a promising direction for future work. Similarly, it is interesting to compare two-player non-stochastic meta-parity games with the, apparently very similar, *hierarchical four player games* (2 vs 2) of [16], which are used to formalize the semantics of first order logic extended with *game quantifiers*. Developing verification methods is another source of possible research directions. For instance, it is interesting to study the *model checking problem* for the logic $\text{pL}\mu^\odot$. Is it possible to compute, or at least approximate to an arbitrary degree of precision, the value $\llbracket F \rrbracket(p)$ assigned by a formula to a state p of a finite PLTS? The problem seems far from trivial. A source of difficulty comes from the fact that (finite) two-player stochastic meta-parity games are not positionally determined. We refer to [25, §8] for an overview of these issues.

Acknowledgments. The author is deeply grateful to his PhD supervisor, Prof. Alex Simpson, who substantially contributed to the research presented in this manuscript with constant insights and suggestions. These thanks are extended to the anonymous reviewers for their valuable comments.

REFERENCES

- [1] F. Bartels. GSOS for probabilistic transition systems. In *Electronic Notes in Theoretical Computer Science, Volume 65, Issue 1*, 2002.
- [2] A. Bianco and L. de Alfaro. Model checking of probabilistic and nondeterministic systems. In *Foundations of Software Technology and Theoretical Computer Science*, number 1026 in Lecture Notes in Computer Science, pages 499–513. Springer-Verlag, 1995.
- [3] D. Blackwell. Infinite G_δ games with imperfect information. *Zastosowania Matematyki (Applicationes Mathematicae)*, Hugo Steinhaus Jubilee Volume X, 1969.
- [4] J. C. Bradfield. Independence: Logics and concurrency. In *In Proceedings of the 14th annual conference on Computer Science Logic (CSL)*, 2000.
- [5] J. C. Bradfield and C. Stirling. Modal logics and mu-calculi: an introduction. In *Handbook of Process Algebra*. Elsevier, 2001.
- [6] T. Brázdil, V. Brozek, A. Kucera, and J. Obdržálek. Qualitative reachability in stochastic BPA games. *Information and Computation*, 209, 2011.
- [7] K. Chatterjee. *Stochastic ω -Regular Games*. PhD thesis, University of California, Berkeley, 2007.
- [8] K. Chatterjee, L. Doyen, and T. A. Henzinger. A survey of stochastic games with $\lim\text{-sup}$ and $\lim\text{-inf}$ objective. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming*. Springer-Verlag Berlin, 2009.

- [9] P. Clairambault, J. Gutierrez, and G. Winskel. The winning ways of concurrent games. In *Logic in Computer Science*. IEEE Computer Society, 2012.
- [10] L. de Alfaro and R. Majumdar. Quantitative solution of omega-regular games. *Journal of Computer and System Sciences, Volume 68, Issue 2*, pages 374 – 397, 2004.
- [11] D. Fischer, E. Grädel, and L. Kaiser. Model checking games for the quantitative μ -calculus. *Theory of Computing Systems*, 47(3), 2010.
- [12] E. Grädel and J. Väänänen. Dependence, independence, and incomplete information. In *International Conference on Database Theory (ICDT)*, 2012.
- [13] O. M. Herescu and C. Palamiessi. Probabilistic asynchronous π -calculus. In *Proceedings of FOSSACS 2000*, pages 146–160, 2000.
- [14] M. Huth and M. Kwiatkowska. Quantitative analysis and model checking. In *Proceedings of the 12th Annual IEEE Symposium on Logic In Computer Science*, page 111, Washington, DC, USA, 1997. IEEE Computer Society.
- [15] T. Jech. *Set Theory*. Springer Monographs in Mathematics. Springer, 2002.
- [16] L. Kaiser. Game quantification on automatic structures and hierarchical model checking games. In *International Conference on Computer Science Logic (CSL)*, pages 411–425, 2006.
- [17] A. S. Kechris. *Classical Descriptive Set Theory*. Graduate Texts in Mathematics. Springer Verlag, 1994.
- [18] D. Kozen. Results on the propositional mu-calculus. In *Theoretical Computer Science*, volume 27, pages 333–354, 1983.
- [19] M. Kwiatkowska, G. Norman, D. Parker, and M. Vigotti. Probabilistic mobile ambients. *Theoretical Computer Science*, pages 12–13, 2009.
- [20] D. A. Martin. Borel determinacy. *Annals of Mathematics*, 2(102):363–371, 1975.
- [21] D. A. Martin. The determinacy of Blackwell games. In *Journal of Symbolic Logic Volume 63, Issue 4*, 1565–1581, 1998.
- [22] D. A. Martin and R. M. Solovay. Internal Cohen extensions. *Ann. Math. Logic*, 2:143–178, 1970.
- [23] A. McIver and C. Morgan. Results on the quantitative μ -calculus $\text{qM}\mu$. *ACM Trans. Comput. Logic*, 8(1):3, 2007.
- [24] M. Mio. Probabilistic Modal μ -Calculus with Independent Product. In *Foundations of Software Science and Computation Structures*, volume 6604 of *Lecture Notes in Computer Science*, pages 290–304. Springer-Verlag Berlin, 2011.
- [25] M. Mio. *Game Semantics for Probabilistic μ -Calculi*. PhD thesis, School of Informatics, University of Edinburgh, 2012. Permanent URL: <http://hdl.handle.net/1842/6223>.
- [26] M. Mio. On the equivalence of denotational and game semantics for the probabilistic μ -calculus. *Logical Methods in Computer Science*, 8(2), 2012.
- [27] C. Morgan and A. McIver. A probabilistic temporal calculus based on expectations. In *In Lindsay Groves and Steve Reeves, editors, Proc. Formal Methods*. Springer Verlag, 1997.
- [28] M. Narasimha, R. Cleaveland, and P. Iyer. Probabilistic temporal logics via the modal mu-calculus. In *Foundations of Software Science and Computation Structures*, pages 288–305, 1999.
- [29] R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, Laboratory for Computer Science, Massachusetts Institute of Technology, 1995.
- [30] C. Stirling. *Modal and temporal logics for processes*. Springer (Texts in Computer Science), 2001.
- [31] T. Tao. *An introduction to measure theory*. Graduate Studies in Mathematics. American Mathematical Society, 2011.
- [32] A. Tarski. A lattice-theoretical fixedpoint theorem and its applications. *Pacific Journal of Mathematics*, 5(2):285–309, 1955.