Theoretical Informatics and Applications Will be set by the publisher
Informatique Théorique et Applications

ON THE EQUIVALENCE OF GAME AND DENOTATIONAL SEMANTICS FOR THE PROBABILISTIC μ-CALCULUS

Matteo Mio[1]

Abstract. The probabilistic modal μ-calculus pLμ (often called the quantitative μ-calculus) is a generalization of the standard modal μ-calculus designed for expressing properties of probabilistic labeled transition systems. The syntax of pLμ formulas coincides with that of the standard modal μ-calculus. Two semantics have been studied for pLμ, both assigning to each process-state p a value in $[0, 1]$ representing the probability that the property expressed by the formula will hold at p: a denotational semantics and a game semantics given by means of two player stochastic (parity) games. The two semantics have been proved to coincide on all finite probabilistic labeled transition systems, but the equivalence of the two semantics on arbitrary models has been open in literature. In this paper we prove that the equivalence indeed holds for arbitrary infinite models, and thus our result strengthens the fruitful connection between denotational and game semantics. Our proof adapts the *unfolding method* recently introduced by Fischer, Grädel and Kaiser.

1991 Mathematics Subject Classification. 91A15, 03B44.

1. Introduction

The modal μ-calculus Lμ [7] is a very expressive logic obtained by extending classical propositional modal logic with least and greatest fixed point operators. The logic Lμ has been extensively studied as it provides a very powerful tool for

Keywords and phrases: probabilistic μ-calculus, stochastic game semantics, determinacy.

[1] LFCS, School of Informatics, University of Edinburgh

expressing properties of labeled transition systems [13]. Encodings of many important temporal logic such as LTL, CTL and CTL* into Lμ [1], provided evidence for the very high expressive power of the calculus. A precise expressivity result was given in [6], where the authors showed that every formula of monadic second order logic over transition systems which does not distinguish between bisimilar models is equivalent to a formula of Lμ. The logic Lμ has a simple denotational interpretation [13]. However it is often very difficult to intuitively grasp the denotational *meaning* of a Lμ formula as the nesting of fixed point operators can induce very complicated properties. To alleviate this problem, another complementary semantics for the logic Lμ, based on two player (parity) games, has been studied in [3, 13]. The two semantics have been proven to coincide and this allows us to pick the most convenient viewpoint when thinking about the logic Lμ.

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) [12] are a natural generalization of standard LTS's to the probabilistic scenario, as they allow both (countable) non-deterministic and probabilistic choices.

The probabilistic modal μ-calculus pLμ, introduced in [2,5,11], is a generalization of Lμ 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 [4], we adopt the *probabilistic* adjective. By this choice, however, we do not mean that the *semantics* of pLμ is *probabilistic*, rather we just want to make clear that the models considered are PLTS's. The syntax of the logic pLμ coincides with that of the standard μ-calculus. The denotational semantics of pLμ [2, 11] generalizes that of Lμ, by interpreting every formula F as a map $[\![F]\!]: P \to [0, 1]$, which assigns to each process p a *degree of truth*. A key aspect of the denotational semantics of [2, 11] is the interpretation of conjunction, defined as $[\![F \wedge G]\!](p) = min\{[\![F]\!](p), [\![G]\!](p)\}$. This is not the only possible meaningful generalization of standard boolean conjunction to the real interval $[0, 1]$; different interpretations for the connectives of pLμ (including the one of [2,11]) have been proposed in [5], and there there is no, *a priori*, good reason to prefer one in favour of the others.

In [10], the authors introduce an alternative semantics for the logic pLμ. This semantics, given in term of two player stochastic (parity) games [14], is a natural generalization of the two player (non stochastic) game semantics for the logic Lμ [13]. As in Lμ 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 Player 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 over process-states (this is the key difference between pLμ and Lμ 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 formulae as expressing, for each process p, the (limit) probability of a *property*, specified by the formula, holding at the state p.

In [10], the equivalence of the denotational and game semantics for pLμ on all *finite* models, was proven. The proof, which adapts the standard technique of [3, 13] used to prove the equivalence of game and denotational semantics for Lμ, makes essential use of the fact that *memoryless* and *optimal* strategies exist in every *finite* two player stochastic game with *parity* objectives [14]. This property however, does not hold, in general, for two player stochastic (parity) games of infinite size: optimal strategies may not exist, and an *unbounded* amount of memory might be necessary even for playing ϵ-optimally, i.e. for guaranteeing a probability of victory ϵ-close to the optimal one. The general result, i.e. the equivalence of the game and denotational semantics for pLμ on arbitrary infinite models, was left open in [10].

In this paper we prove that the equivalence indeed holds for arbitrary infinite models, thus strengthening the connection between denotational and game semantics. This result provides an *a posteriori* justification for the denotational interpretation in [2, 11] of the connectives of pLμ. Moreover, the extension of the result of [10] to arbitrary infinite models is of practical interest, since infinite state systems often provide natural abstractions for, e.g., infinite memory, infinite data-sets *etcetera*. Our contribution consists in adapting a technique recently introduced in [4], called the *unfolding method*, which the authors used to prove a similar result for a μ-calculus-like logic designed to express quantitative properties of (non probabilistic) labeled transition systems. While this is not a difficult adaption, the result seems worth notice since the question has been open in literature since [10]. Moreover the differences between the games considered in [4] and pLμ stochastic games, e.g. the fact that *Markov chains* are the outcomes of the games rather than just infinite paths, make this result not immediate from [4].

The rest of the paper is organized as follows: in section 2, we introduce some mathematical definitions; in section 3, we define the syntax and the denotational semantics of the logic pLμ as in [2, 11]; in section 4, we define the class of parity games that are going to be used to give game semantics to the logic; in section 5, we define the game semantics of pLμ in terms of two player stochastic parity games as in [2, 10] and state the main theorem which asserts the equivalence of the denotational and game semantics for pLμ; lastly, in section 6, a detailed proof of the main theorem is given.

2. Background definitions and notation

Definition 2.1 (Probability distributions)**.** A probability distribution d over a set X is a function $d: X \to [0,1]$ such that $\sum_{x \in X} d(x) = 1$. The *support* of d, denoted by $supp(d)$ is defined as the 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 distribution over X such that $supp(\delta_x) = \{x\}$, i.e. the unique distribution such that $\delta_x(x) = 1$ and $\delta_x(y) = 0$, for all $y \neq x$.

Definition 2.2 (PLTS [12])**.** Given a countable set L of labels, a *probabilistic labeled transition system* is a pair $\langle P, \{\xrightarrow{a}\}_{a \in L}\rangle$, where P is a *countable* set of process-states and $\xrightarrow{a} \subseteq P \times \mathcal{D}(P)$ for every $a \in L$. As usual we write $p \xrightarrow{a} d$ if $(p,d) \in \xrightarrow{a}$. We refer to the set $\bigcup_{p \in P} \bigcup_{a \in L} \{d \mid p \xrightarrow{a} d\}$, as the *set of distributions of the PLTS*. We say that a PLTS is *countably branching* if its set of distributions is countable. In what follows, we restrict our attention to countably branching PLTS's only.

The transition relation of a PLTS models the dynamics of the processes: $p \xrightarrow{a} d$ means that the process p can perform the atomic action $a \in L$ and then, with probability $d(q)$, behaves like the process q. Probabilistic labeled transition system are a natural generalization of labeled transition systems to the probabilistic scenario: a standard LTS can be modeled as a PLTS, in which every reachable distribution is of the form δ_p, for some $p \in P$.

Definition 2.3 (Lattice operations)**.** Given a set X, we denote with 2^X the set of all sets $Y \subseteq X$. Given a complete lattice $(X, <)$, we denote with $\bigsqcup : 2^X \to X$ and $\bigsqcap : 2^X \to X$ the operations of join and meet respectively.

In the following, we assume standard notions of basic topology and basic measure theory. The topological spaces we consider will always be 0-dimensional Polish spaces.

3. The Probabilistic Modal μ-Calculus

Given a set Var of propositional variables ranged over by the letters X, Y, Z and a set of labels L ranged over by the letters a, b, c, the formulae of the logic are defined by the following grammar:

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

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 (Subformulae)**.** We define the set $Sub(F)$ of *subformualae* of F by case analysis on F as follows:

$$
\begin{array}{lcl}
Sub(X) & \stackrel{\text{def}}{=} & \{X\} \\
Sub(F_1 \wedge F_2) & \stackrel{\text{def}}{=} & \{F_1 \wedge F_2\} \cup Sub(F_1) \cup Sub(F_2) \\
Sub(F_1 \vee F_2) & \stackrel{\text{def}}{=} & \{F_1 \vee F_2\} \cup Sub(F_1) \cup Sub(F_2) \\
Sub([a]\, F_1) & \stackrel{\text{def}}{=} & \{[a]\, F_1\} \cup Sub(F_1) \\
Sub(\langle a \rangle F_1) & \stackrel{\text{def}}{=} & \{\langle a \rangle F_1\} \cup Sub(F_1) \\
Sub(\nu X.F_1) & \stackrel{\text{def}}{=} & \{\nu X.F_1\} \cup Sub(F_1) \\
Sub(\mu X.F_1) & \stackrel{\text{def}}{=} & \{\nu X.F_1\} \cup Sub(F_1)
\end{array}
$$

We say that G is a subformula of F if $G \in Sub(F)$.

Definition 3.2 (Normal Formula)**.** We say that a 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.

For convenience we only consider, from now on, formulae F in normal form. For example this allows the definition below to be given as follows:

Definition 3.3 (Variable subsumption)**.** Given a formula F, we say that X *subsumes* Y in F, for $X \neq Y$, if X and Y are bound in F by the sub-formulae $\star_1 X.G$ and $\star_2 Y.H$ respectively, and $\star_2 Y.H \in Sub(G)$, for $\star_1, \star_2 \in \{\mu, \nu\}$.

Given a PLTS $\langle P, \{\xrightarrow{a}\}_{a \in L}\rangle$, we denote with $(P \to [0,1])$ and with $(\mathcal{D}(P) \to [0,1])$ the complete lattices of functions from P and from $\mathcal{D}(P)$ respectively, to the real interval $[0,1]$ with the pointwise order. Given a function $f : P \to [0,1]$, we denote with $\overline{f} : \mathcal{D}(P) \to [0,1]$ the lifted function defined as follows:

$$\overline{f} \stackrel{\text{def}}{=} \lambda d. \left(\sum_{p \in P} d(p) \cdot f(p) \right)$$

A function $\rho : Var \to (P \to [0,1])$ is called an *interpretation* of the variables. Given a function $f : P \to [0,1]$ 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 $[\![F]\!]_\rho : P \to [0,1]$ of the pLμ formula F under the interpretation ρ, is defined by structural induction on F as follows:

$$
\begin{array}{lcl}
[\![X]\!]_\rho & = & \rho(X) \\
[\![G \vee H]\!]_\rho & = & [\![G]\!]_\rho \sqcup [\![H]\!]_\rho \\
[\![G \wedge H]\!]_\rho & = & [\![G]\!]_\rho \sqcap [\![H]\!]_\rho \\
[\![\langle a \rangle G]\!]_\rho & = & \lambda p. \left(\bigsqcup \{ \overline{[\![G]\!]_\rho}(d) \mid p \xrightarrow{a} d \} \right) \\
[\![[a]\, G]\!]_\rho & = & \lambda p. \left(\bigsqcap \{ \overline{[\![G]\!]_\rho}(d) \mid p \xrightarrow{a} d \} \right) \\
[\![\mu X.G]\!]_\rho & = & \text{least fixed point of } \lambda f.([\![G]\!]_{\rho[f/X]}) \\
[\![\nu X.G]\!]_\rho & = & \text{greatest fixed point of } \lambda f.([\![G]\!]_{\rho[f/X]})
\end{array}
$$

Since the interpretation assigned to every pLμ operator is monotone, the existence of the least and greatest fixed points is guaranteed by the Knaster-Tarski theorem.

Moreover the least and the greatest fixed points can be computed inductively as:

$$[\![\mu X.G]\!]_\rho = \bigsqcup_\alpha [\![\mu X.G]\!]^\alpha_\rho \quad \text{and} \quad [\![\nu X.G]\!]_\rho = \bigsqcap_\alpha [\![\nu X.G]\!]^\alpha_\rho$$

where $[\![\mu X.G]\!]^\alpha_\rho$ and $[\![\nu X.G]\!]^\alpha_\rho$ are defined as:

$$[\![\mu X.G]\!]^\alpha_\rho = \bigsqcup_{\beta<\alpha} [\![G]\!]_{\rho[[\![\mu X.G]\!]^\beta_\rho/X]} \quad \text{and} \quad [\![\nu X.G]\!]^\alpha_\rho = \bigsqcap_{\beta<\alpha} [\![G]\!]_{\rho[[\![\nu X.G]\!]^\beta_\rho/X]}$$

with α and β ranging over the ordinals.

4. Two Player Stochastic Parity Games

In this section we introduce the class two player stochastic games used to give game semantics to the logic pLμ. This material in standard, and follows similar presentations, as in e.g. [14].

A two player turn-based stochastic game (or just a $2\frac{1}{2}$-player game) is played on some *arena* $\mathcal{A} = \langle (S,E),(S_1,S_2,S_N),\pi\rangle$ where (S,E) is a directed graph with *countable* set of states S and transition relation $E \subseteq S\times S$; the sets S_1, S_2, S_N are a partition of S and $\pi : S_N \to \mathcal{D}(S)$ is called the probabilistic transition function. For every state $s\in S$, we denote with $E(s)$ the (possibly infinite) set $\{s' \mid (s,s')\in E\}$ of E-successors of s. We require that for all $s\in S_N$, $E(s) = supp(\pi(s))$. We denote with S_t the set of *terminal* states, i.e. those $s\in S$ such that $E(s) = \emptyset$.

The game is played on the arena $\mathcal{A}$ by three players named Player 1, Player 2 and Nature, the probabilistic agent of the game. The states in S_1 are under the control of Player 1, the states in S_2 are under the control of Player 2, and the states in S_N are probabilistic, i.e. under the control of Nature. At a state $s\in S_1$, if $s \notin S_t$, Player 1 chooses a successor from the set $E(s)$; if $s\in S_t$ the game ends. Similarly, at a state $s\in S_2$, if $s\notin S_t$, Player 2 chooses a successor from the set $E(s)$; if $s\in S_t$ the game ends. At a state $s\in S_N$, a successor state is probabilistically chosen by Nature, according to the distribution $\pi(s)$. The outcome of a *play* of the three players is a path in $\mathcal{A}$, either infinite or finite (ending in a terminal state), which we call a *completed* path.

Definition 4.1. We denote with $\mathcal{P}^\omega$ and $\mathcal{P}^{<\omega}$ the sets of infinite and finite (non empty) paths in $\mathcal{A}$. Given a finite path $\vec{s}\in\mathcal{P}^{<\omega}$ we denote with $last(\vec{s})$ the last state $s\in S$ of $\vec{s}$. We write $\vec{s} \lhd \vec{t}$, with $\vec{s},\vec{t}\in\mathcal{P}^{<\omega}$, if $\vec{t} = \vec{s}.s$, for some $s\in S$, where as usual the *dot* symbol denotes the concatenation operator; the relation $\lhd$ is called the *successor relation on finite 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 $last(\vec{s}) \in S_t$. We denote with $\mathcal{P}^{<\omega}_1$, $\mathcal{P}^{<\omega}_2$ and $\mathcal{P}^{<\omega}_N$ the sets of finite paths $\vec{s}$ such that $last(\vec{s})\in S_1$, $last(\vec{s})\in S_2$ and $last(\vec{s})\in S_N$ 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. We denote with $(\mathcal{P},\Omega)$ the Borel σ-algebra induced by the topology on $\mathcal{P}$, i.e. the smallest σ-algebra on $\mathcal{P}$ containing all the open sets.

To specify the reward assigned to Player 1 when a given completed path $\vec{s}$ is the outcome of a play, we introduce the notion of payoff function.

Definition 4.2. A *payoff function* for the arena $\mathcal{A}$ is a Borel-measurable function $\Phi:\mathcal{P}\rightarrow[0,1]$. The value $\Phi(\vec{s})$, for a given $\vec{s}\in\mathcal{P}$, should be understood as the reward assigned to Player 1 when $\vec{s}$ is the outcome of a play in $\mathcal{A}$.

Definition 4.3 (Two player stochastic game)**.** A *two player turn based stochastic game* (or just a $2\frac{1}{2}$-player game) is a pair $\langle\mathcal{A},\Phi\rangle$, where Φ is a payoff function for the arena $\mathcal{A}$.

The goal of Player 1 in the game $\langle\mathcal{A},\Phi\rangle$ is to maximize his payoff, while the *dual* goal of Player 2 is to minimize the payoff assigned to Player 1.

When working with *stochastic* games, it is useful to look at the possible outcomes of a play up-to the behavior of Nature. This is done by introducting the notion of Markov chain in $\mathcal{A}$, whose precise formulation is provided by the following definitions.

Definition 4.4 (Tree in $\mathcal{A}$)**.** A *tree* in the arena $\mathcal{A}$ is a collection $T\subseteq\mathcal{P}^{<\omega}$ of finite paths in $\mathcal{A}$, such that

(1) T is down-closed: if $\vec{s}\in T$ and $\vec{t}$ is a prefix of $\vec{s}$, then $\vec{t}\in T$.

(2) T has a root: there exists exactly one finite path $\vec{s}=(s_0)$ of length one in T. The path $\vec{s}$, denoted by $root(T)$, is called the root of the tree T.

The set of *children* of the node $\vec{s}$ in T is the set $\{\vec{t}\in T \mid \vec{s}\lhd\vec{t}\}$. We consider the nodes $\vec{s}$ of T as labeled by the *last* function.

Definition 4.5 (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 $last(\vec{s})\in S_t$ or $\vec{s}$ has a unique child in T. Similarly, $\vec{s}$ is *fully branching* in T if, for every $s\in E(last(\vec{s}))$, it holds that $\vec{s}.s\in T$.

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

(1) If $last(\vec{s})\in S_1\cup S_2$ then $\vec{s}$ branches uniquely in M.

(2) If $last(\vec{s})\in S_N$ then $\vec{s}$ branches fully in M.

Definition 4.7 (Probability measure $\mathcal{M}$)**.** Every Markov chain M determines a probability assignment $\mathcal{M}(O_{\vec{s}})$ to every basic clopen set $O_{\vec{s}}\subseteq\mathcal{P}$, for $\vec{s}$ a finite path $\vec{s}=(s_0,s_1,...,s_n)$ with $n\in\mathbb{N}$, defined as follows:

$$\mathcal{M}(O_{\vec{s}}) \stackrel{\text{def}}{=} \begin{cases} \prod\{\pi(s_i)(s_{i+1}) \mid 0 \leq i < n \ \wedge \ s_i \in S_N\} & \text{if } \vec{s} \in M \\ 0 & \text{otherwise} \end{cases}$$

In other words, $\mathcal{M}$ assigns to the basic open set $O_{\vec{s}} \subseteq \mathcal{P}$, i.e the set of all completed paths having $\vec{s}$ as prefix, value 0 if $\vec{s}$ is not a path in M, and the product of all probabilities labeling the probabilistic steps in $\vec{s}$, otherwise; note that if there are no probabilistic steps in $\vec{s}$, $\mathcal{M}$ assigns to $O_{\vec{s}}$ probability 1, which is the value of the empty product. The assignment $\mathcal{M}$ extends to a unique probability measure on the Borel σ-algebra $(\mathcal{P}, \Omega)$, which is the smallest σ-algebra on $\mathcal{P}$ containing all the open sets [14].

Given the previous definitions we can define the *expected reward* of Player 1 when a given Markov chain M is the result (up to the behavior of Nature) of a play.

Definition 4.8 (Expected reward of M)**.** Let $\langle \mathcal{A}, \Phi \rangle$ be a $2\frac{1}{2}$-player game. We define the *expected reward* of a Markov chain M in $\mathcal{A}$, denoted by $E(M)$, as follows:

$$E(M) = \int_{\mathcal{P}} \Phi \ d\mathcal{M}_M$$

where $\mathcal{M}_M$ denotes the probability measure on the Borel σ-algebra $(\mathcal{P}, \Omega)$, induced by M, as described above.

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

Definition 4.9. An *unbounded memory deterministic strategy* (or just a strategy) σ_1 for Player 1 in $\mathcal{A}$ is defined as a function $\sigma_1 : \mathcal{P}_1^{<\omega} \to S \cup \{\bullet\}$ such that $\sigma_1(\vec{s}) \in E(last(s))$ if $E(last(\vec{s})) \neq \emptyset$ and $\sigma_1(\vec{s}) = \bullet$ otherwise. Similarly a strategy σ_2 for Player 2 is defined as a function $\sigma_2 : \mathcal{P}_2^{<\omega} \to S \cup \{\bullet\}$. We say that a strategy σ_1 for Player 1 is *memory-less*, if there exists a function $f : S_1 \to S \cup \{\bullet\}$ such that for every $\vec{s} \in \mathcal{P}_1^{<\omega}$, the equality $\sigma_1(\vec{s}) = f(last(\vec{s}))$ holds. Similarly, a strategy σ_2 for Player 2 is memory-less if there exists a function $f : S_2 \to S \cup \{\bullet\}$ such that for every $\vec{s} \in \mathcal{P}_2^{<\omega}$, the equality $\sigma_2(\vec{s}) = f(last(\vec{s}))$ holds. In other words a strategy is memory-less if its decision on any history $\vec{s}$, only depends on the last state $last(\vec{s})$ of $\vec{s}$. 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.

Definition 4.10 ($M^{s_0}_{\sigma_1,\sigma_2}$)**.** Given an initial state $s_0 \in S$ and a strategy profile $\langle \sigma_1, \sigma_2 \rangle$, a unique Markov chain $M^{s_0}_{\sigma_1,\sigma_2}$ is determined:

(1) the root of $M^{s_0}_{\sigma_1,\sigma_2}$ is labeled with s_0,

(2) for every $\vec{s} \in M^{s_0}_{\sigma_1,\sigma_2}$, if $last(\vec{s}) = s$ with $s \in S_1$ not a terminal state, then the unique child of $\vec{s}$ in $M^{s_0}_{\sigma_1,\sigma_2}$ is $\vec{s}.\big(\sigma_1(\vec{s})\big)$,

(3) for every $\vec{s} \in M^{s_0}_{\sigma_1,\sigma_2}$, if $last(\vec{s}) = s$ with $s \in S_2$ not a terminal state, then the unique child of $\vec{s}$ in $M^{s_0}_{\sigma_1,\sigma_2}$ is $\vec{s}.\big(\sigma_2(\vec{s})\big)$.

We denote with $\mathcal{M}^{s_0}_{\sigma_1,\sigma_2}$ the probability measure over $\langle \mathcal{P}, \Omega \rangle$ associated with the Markov chain $M^{s_0}_{\sigma_1,\sigma_2}$.

Definition 4.11. Given a $2\frac{1}{2}$-player game $\mathcal{G}=\langle\mathcal{A},\Phi\rangle$ and an initial state $s\in S$, we define the *lower value* and *upper value* of the game $\mathcal{G}$ at s, denoted by $Val_{\uparrow}(\mathcal{G})(s)$ and $Val_{\downarrow}(\mathcal{G})(s)$ respectively, as follows:

$$Val_{\downarrow}(\mathcal{G})(s) = \bigsqcup_{\sigma_1}\bigsqcap_{\sigma_2} E(M^s_{\sigma_1,\sigma_2}) \qquad Val_{\uparrow}(\mathcal{G})(s) = \bigsqcap_{\sigma_2}\bigsqcup_{\sigma_1} E(M^s_{\sigma_1,\sigma_2}).$$

$Val_{\downarrow}(\mathcal{G})(s)$ represents the (limit) expected reward that Player 1 can get, 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 $Val_2(\mathcal{G})(s)$ represents the (limit) expected reward that Player 1 can get, 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 $Val_{\downarrow}(\mathcal{G})(s) \leq Val_{\uparrow}(\mathcal{G})(s)$ for every $s\in S$.

Theorem 4.12 (Determinacy [8, 9])**.** *For every $2\frac{1}{2}$-player game $\mathcal{G}=\langle\mathcal{A},\Phi\rangle$ the following equality holds:*

$$\forall s\in S,\ Val_{\downarrow}(\mathcal{G})(s) = Val_{\uparrow}(\mathcal{G})(s).$$

Intuitively the determinacy theorem states that the players do not get any advantage by letting the opponent choose his strategy first. We just write $Val(\mathcal{G})(s)$ for the *value* of the game at s defined as $Val_{\downarrow}(\mathcal{G})(s) = Val_{\uparrow}(\mathcal{G})(s)$.

Definition 4.13 (ϵ-optimal strategies)**.** Given a $2\frac{1}{2}$-player game $\mathcal{G}=\langle\mathcal{A},\Phi\rangle$, a strategy σ_1 for Player 1 is called ϵ-*optimal* at s, for some $\epsilon\in[0,1]$, if the following inequality holds:

$$\bigsqcap_{\sigma_2} E(M^s_{\sigma_1,\sigma_2}) > Val(\mathcal{G})(s) - \epsilon.$$

Similarly a strategy σ_2 for Player 2 is called ϵ-*optimal* at s, if the following inequality holds:

$$\bigsqcup_{\sigma_1} E(M^s_{\sigma_1,\sigma_2}) < Val(\mathcal{G})(s) + \epsilon.$$

We refer to a strategy as *optimal* as s, if it is 0-optimal as s. Clearly, from theorem 4.12, for every $\epsilon>0$ and for every $s\in S$ there exist ϵ-optimal strategies for Player 1 and Player 2 at s. However, in general, there might be no 0-optimal strategies [14].

In this paper we are interested in $2\frac{1}{2}$-player *parity* games, which are those $2\frac{1}{2}$-player games $\langle\mathcal{A},\Phi\rangle$ where the payoff function Φ is induced by a *parity structure* $\langle\mathbb{P},\mathbb{B}\rangle$.

Definition 4.14. Given a $2\frac{1}{2}$-player arena $\mathcal{A}$, a *parity structure* for $\mathcal{A}$ is a pair $\langle\mathbb{P},\mathbb{B}\rangle$, where $\mathbb{P}$ (called the *priority assignment*) is a function $\mathbb{P}: S\rightarrow\{0,...,n\}$, for some $n\in\mathbb{N}$, assigning a natural number in $\{0,...,n\}$ to every state in S, and $\mathbb{B}$ (called the *terminal reward assignment*) is a function $\mathbb{B}: S_t\rightarrow[0,1]$ assigning a value in the real interval $[0,1]$ to each terminal state $s\in S_t$. Given a priority assignment $\mathbb{P}$, and an infinite path $\vec{s}=(s_i)_{i\in\mathbb{N}}$ in $\mathcal{P}^{\omega}$, the smallest natural number appearing infinitely often in the infinite sequence $(\mathbb{P}(s_i))_{i\in\mathbb{N}}$, is denoted, with a little abuse of notation, by $\mathbb{P}(\vec{s})$.

A parity structure $\langle \mathbb{P}, \mathbb{B} \rangle$ for the arena $\mathcal{A}$, determines a unique payoff function $\Phi_{\langle \mathbb{P},\mathbb{B}\rangle} : \mathcal{P} \rightarrow [0,1]$ defined as follows:

$$\Phi_{\langle \mathbb{P},\mathbb{B}\rangle}(\vec{s}) = \begin{cases} \mathbb{B}(last(\vec{s})) & \text{if } \vec{s} \in \mathcal{P}^t \\ 1 & \text{if } \vec{s} \in \mathcal{P}^\omega \text{ and } \mathbb{P}(\vec{s}) \text{ is even} \\ 0 & \text{if } \vec{s} \in \mathcal{P}^\omega \text{ and } \mathbb{P}(\vec{s}) \text{ is odd} \end{cases}$$

The function $\Phi_{\langle \mathbb{P},\mathbb{B}\rangle}$ is Borel-measurable for every parity structure $\langle \mathbb{P}, \mathbb{B} \rangle$ [14]. Note that by the above definition, we do not impose any constraint on the reward assigned by $\mathbb{B}$ on the terminal states in $S_t \cap S_1$ and in $S_t \cap S_2$. However it will be useful later, when considering pLμ-games, to impose $\mathbb{B}(s) = 0$ for all $s \in S_t \cap S_1$, and $\mathbb{B}(s) = 1$ for all $s \in S_t \cap S_2$, to implement that a player loses if they get stuck.

A $2\frac{1}{2}$-player *parity* game is thus defined as a $2\frac{1}{2}$-player game $\mathcal{G} = \langle \mathcal{A}, \Phi \rangle$, where $\Phi = \Phi_{\langle \mathbb{P},\mathbb{B}\rangle}$, for some priority structure $\langle \mathbb{P}, \mathbb{B} \rangle$ on $\mathcal{A}$.

We now state an important property (often called *prefix independence*) of parity payoff functions:

Lemma 4.15. *Let $\mathcal{G} = \langle \mathcal{A}, \Phi \rangle$ be $2\frac{1}{2}$-player parity game. Then for every completed paths $\vec{s}$ and $\vec{t}$ in $\mathcal{A}$, such that $\vec{s} = \vec{s'}.\vec{t}$, for some finite path $\vec{s'}$ in $\mathcal{A}$, then the equality $\Phi(\vec{s}) = \Phi(\vec{t})$ holds.*

Proposition 4.16 ([14])**.** *There exists a $2\frac{1}{2}$-player parity game (with infinite state space S), such that no 0-optimal strategy exists for either player. Moreover, for any $\epsilon > 0$, no memory-less ϵ-optimal strategy exists for either player.*

5. Game Semantics

In this section we define the *game semantics* of the probabilistic modal μ-calculus, in terms of $2\frac{1}{2}$-player parity games.

Given a PLTS $\langle P, \{\xrightarrow{a}\}_{a\in L}\rangle$ (with countable set of distributions D, as defined in definition 2.2), a pLμ formula F and an interpretation ρ of the variables, we denote with $\mathcal{G}^F_\rho$ the parity game $\langle A, \Phi_{\mathbb{P},\mathbb{B}}\rangle$ formally defined as follows. The state space of the arena $\mathcal{A} = \langle (S, E), \{S_1, S_2, S_N\}, \pi \rangle$, is the set $S = (P \times Sub(F)) \cup (D \times Sub(F))$ of pairs of states $p \in P$ or distributions $d \in D$, and subformulae $G \in Sub(F)$; the transition relation E is defined as $E(\langle d, G\rangle) = \{\langle p, G\rangle \mid p \in supp(d)\}$ for every $d \in D$; $E(\langle p, G\rangle)$ is defined by case analysis on the outermost connective of G as follows:

(1) if $G = X$, with X free in F, then $E(\langle p, G\rangle) = \emptyset$.

(2) if $G = X$, with X bound in F by the subformula $\star X.H$, with $\star \in \{\mu, \nu\}$, then $E(\langle p, G\rangle) = \{\langle p, H\rangle\}$.

(3) if $G = \star X.H$, with $\star \in \{\mu, \nu\}$, then $E(\langle p, G\rangle) = \{\langle p, H\rangle\}$.

(4) if $G = \langle a \rangle H$ or $G = [a]\, H$ then $E(\langle p, G\rangle) = \{\langle d, H\rangle \mid p \xrightarrow{a} d\}$.

(5) if $G = H \vee H'$ or $G = H \wedge H'$ then $E(\langle p, G\rangle) = \{\langle p, H\rangle, \langle p, H'\rangle\}$

The partition $\{S_1, S_2, S_N\}$ 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 $\langle d, G\rangle$ is in S_N. Finally, the terminal states $\langle p, X\rangle$, with X free in F, are in S_1 by convention. The probability transition function $\pi : S_N \to \mathcal{D}(S)$ is defined as $\pi(\langle d, G\rangle)(\langle p, G\rangle) = d(p)$. The priority assignment $\mathbb{P}$ is defined as usual in μ-calculus model checking games (see e.g. [4]): the priority assigned to the states $\langle p, X\rangle$, with X a μ-variable is odd; dually the priority assigned to the states $\langle p, X\rangle$, with X a ν-variable is even. Moreover $\mathbb{P}(\langle p, X\rangle) < \mathbb{P}(\langle p', X'\rangle)$ if X subsumes X' in F. All other states get priority n, for some big enough $n \in \mathbb{N}$ such that $n > \mathbb{P}(\langle p, X\rangle)$, for every bound variable X in F. The terminal reward assigment $\mathbb{B}$ is defined as $\mathbb{B}(\langle p, X\rangle) = \rho(X)(p)$ for every terminal state $\langle p, X\rangle$ with X free in F. All other terminal states in $\mathcal{G}^F_\rho$ are either of the form $\langle p, \langle a\rangle H\rangle$ or $\langle p, [a]\, H\rangle$, with $\{d \mid p \xrightarrow{a} d\} = \emptyset$. The reward assignment $\mathbb{B}$ is defined on these terminal states as follows: $\mathbb{B}(\langle p, \langle a\rangle H\rangle) = 0$ and $\mathbb{B}(\langle p, [a]\, H\rangle) = 1$. As noted before, this implements the policy that a player loses if they get stuck.

Definition 5.1. Fix a PLTS $\langle P, \{\xrightarrow{a}\}_{a\in L}\rangle$. The *game semantics* of the formula F under the interpretation ρ, is the function $(\!| F |\!)_\rho : P \to [0,1]$ defined as $(\!| F |\!)_\rho = \lambda p.Val(\mathcal{G}^F_\rho)(\langle p, F\rangle)$, i.e. the function that maps each process-state p to the *value* of the game $\mathcal{G}^F_\rho$ at $\langle p, F\rangle$.

We are now ready to state our main theorem which asserts the equivalence of the denotational and game semantics of the logic pLμ on all PLTS's.

Theorem 5.2. *Given an arbitrary PLTS $\langle P, \{\xrightarrow{a}\}_{a\in L}\rangle$, for every pL$\mu$ formula F and interpretation ρ for the variables, the following equality holds: $[\![F]\!]_\rho = (\!| F |\!)_\rho$.*

The proof of theorem 5.2 is carried out in full detail in the following section.

6. Proof of theorem 5.2

As anticipated in the introduction, the main difficulty in proving theorem 5.2 is that in general, as stated in lemma 4.16, 0-optimal strategies, or even memoryless ϵ-optimal strategies may not exist in a given pLμ game. This compels us to use a different technique than the ones used in, e.g. [10,13], which require the existence of optimal memoryless strategies.

The proof technique we adopt is similar to the *unfolding method* of [4]. 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 2 priorities; then the the general result for games of $n+1$ priority follows by some argument making use of the inductive hypothesis.

In our setting we are interesting in pLμ games of the form $\mathcal{G}^F_\rho$, and the property we want to prove is that the value 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 $\mathcal{G}^F_\rho$. This allows a more transparent and arguably simpler proof.

What we actually prove below, is a stronger result which implies theorem 5.2 and also proves that pLμ games are *determined*. Even though the determinacy of pLμ games follows from theorem 4.12 [9], it is arguably useful to provide a purely inductive proof of determinacy for the simple class (compared to wide class of Blackwell games of [9]) of pLμ games.

Theorem 6.1. *Given an arbitrary PLTS* $\langle P, \{\xrightarrow{a}\}_{a\in L}\rangle$*, for every pL$\mu$ formula* F*, interpretation* ρ *and process-state* $p\in P$*, the following inequalities hold:*

(1) $[\![F]\!]_\rho(p) \leq Val_\downarrow(\mathcal{G}^F_\rho)(\langle p, F\rangle)$

(2) $[\![F]\!]_\rho(p) \geq Val_\uparrow(\mathcal{G}^F_\rho)(\langle p, F\rangle)$

which in turn, imply the following equalities:

(3) $[\![F]\!]_\rho(p) = Val_\downarrow(\mathcal{G}^F_\rho)(\langle p, F\rangle) = Val_\uparrow(\mathcal{G}^F_\rho)(\langle p, F\rangle) = (\!| F |\!)_\rho(p)$.

Proof. Let us a fix an arbitrary PLTS $\langle P, \{\xrightarrow{a}\}_{a\in L}\rangle$. The proof proceeds by induction on the structure of F. We just prove the cases for $F = X$, $F = G \vee H$, $F = \langle a\rangle G$ and $F = \mu X.G$, as the proofs for the dual operators are similar. The most difficult step is in the proof on the inequality 2 in the case $F = \mu X.G$.

Case $\mathbf{F} = \mathbf{X}$, for some variable $X \in Var$. For every process-state $p \in P$ and every interpretation ρ, we have by definition $[\![X]\!]_\rho(p) = \rho(X)(p)$. In the game $\mathcal{G}^X_\rho$ the state $\langle p, X\rangle$ is *terminal* (and therefore the game immediately terminates when starting at this state) and the *terminal reward assignment* $\mathbb{B}$ of $\mathcal{G}^X_\rho$, assigns a payoff of $\rho(X)(p)$ to Player 1. Therefore the lower and upper values at $\langle p, X\rangle$ of the (trivial) game $\mathcal{G}^X_\rho$ are $Val_\downarrow(\mathcal{G}^X_\rho)(\langle p, X\rangle) = Val_\uparrow(\mathcal{G}^X_\rho)(\langle p, X\rangle) = \rho(X)(p)$ by definition.

Case $\mathbf{F} = \mathbf{G} \vee \mathbf{H}$. For every process-state $p \in P$ and every interpretation ρ, we have by definition $[\![G \vee F]\!]_\rho(p) = [\![G]\!]_\rho(p) \sqcup [\![H]\!]_\rho(p)$. Let us consider the state $\langle p, G \vee H\rangle$ of the game $\mathcal{G}^{G\vee H}_\rho$. This state is in S_1, i.e. under the control of Player 1, which can choose to move either to $\langle p, G\rangle$ or $\langle p, H\rangle$. Moreover observe that once the state $\langle p, G \vee H\rangle$ is left after the initial move, it is not reachable again in the game; and that once reached the state $\langle p, G\rangle$ the rest of the game is identical to the game $\mathcal{G}^G_\rho$ (starting at $\langle p, G\rangle$), and similarly once reached the state $\langle p, H\rangle$, the rest of the game is identical to the game $\mathcal{G}^H_\rho$ (starting at $\langle p, H\rangle$).

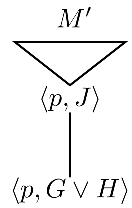

FIGURE 1. Markov chain $M^{\langle p,G\vee H\rangle}_{\sigma_1^\epsilon,\sigma_2}$ in $\mathcal{G}^{G\vee H}_\rho$

Let us first consider the inequality $[\![F]\!]_\rho(p)\leq Val_\downarrow(\mathcal{G}^F_\rho)(\langle p,F\rangle)$, or equivalently by definition 4.11, $[\![F]\!]_\rho(p)\leq\bigsqcup_{\sigma_1}\bigsqcap_{\sigma_2}E(M^{\langle p,F\rangle}_{\sigma_1,\sigma_2})$. We show that this inequality holds, by showing that for each $\epsilon>0$, the inequality $[\![F]\!]_\rho(p)-\epsilon\ <\bigsqcup_{\sigma_1}\bigsqcap_{\sigma_2}E(M^{\langle p,F\rangle}_{\sigma_1,\sigma_2})$ holds. This is done by constructing a strategy σ_1^ϵ, for Player 1 in the pLμ game $\mathcal{G}^F_\rho$, satisfying the following inequality: $[\![F]\!]_\rho(p)-\epsilon\ <\bigsqcap_{\sigma_2}E(M^{\langle p,F\rangle}_{\sigma_1^\epsilon,\sigma_2})$. Let us denote with J the disjunct subformula of F that maximizes $[\![\_]\!]_\rho$ at p, i.e. the formula $J\in\{G,H\}$ such that $[\![J]\!]_\rho(p)=[\![G\vee H]\!]_\rho(p)$. By induction hypothesis, we know that $[\![J]\!]_\rho(p)=Val_\downarrow(\mathcal{G}^J_\rho)(\langle p,J\rangle)$. Therefore, for every $\epsilon>0$, there exists a strategy τ_1^ϵ for Player 1 in the pLμ game $\mathcal{G}^J_\rho$ such that $[\![J]\!]_\rho(p)-\epsilon<\bigsqcap_{\tau_2}E(M^{\langle p,J\rangle}_{\tau_1^\epsilon,\tau_2})$. The strategy σ_1^ϵ for Player 1 in the pLμ game $\mathcal{G}^{G\vee H}_\rho$ is defined as follows:

$$\sigma_1^\epsilon(\vec{s})\stackrel{\text{def}}{=}\begin{cases}\langle q,J\rangle & \text{if } \vec{s}=(\langle q,G\vee H\rangle) \text{ for } q\in P\\ \tau_1^\epsilon(\vec{t}) & \text{if } \vec{s}=\langle q,G\vee H\rangle.\vec{t} \text{ for } q\in P\\ ? & \text{otherwise}\end{cases}$$

The strategy σ_1^ϵ can be described informally as follows: at the initial states of the form $\langle q,G\vee H\rangle$, with $q\in P$ (and in particular on the initial state $\langle p,G\vee H\rangle$), Player 1 chooses to move to $\langle q,J\rangle$. After this choice, his next decisions, based on some play-history $\vec{s}=\langle p,G\vee H\rangle.\vec{t}$, will coincide with the decision taken by τ_1^ϵ on the history $\vec{t}$. Note that $\vec{t}$ is indeed a finite path in the game $\mathcal{G}^J_\rho$ starting in $\langle p,J\rangle$ as observed above, and therefore σ_1^ϵ is well defined on all game-histories starting at $\langle q,F\vee G\rangle$. The question mark in the last clause in definition of σ_1^ϵ represent an unspecified behavior of σ_1^ϵ on all game-histories starting in a state not of the form $\langle q,G\vee H\rangle$, for $q\in P$; a precise specification of σ_1^ϵ can be obtained by substituting the question mark symbol with $\sigma_1'(\vec{s})$, for whatever strategy $\sigma_1'(\vec{s})$ for Player 1 in the game $\mathcal{G}^{G\vee H}_\rho$. We now show that the inequality $[\![G\vee H]\!]_\rho(p)-\epsilon\ <\bigsqcap_{\sigma_2}E(M^{\langle p,G\vee H\rangle}_{\sigma_1^\epsilon,\sigma_2})$ holds, or equivalently by definition of J, $[\![J]\!]_\rho(p)-\epsilon\ <\bigsqcap_{\sigma_2}E(M^{\langle p,G\vee H\rangle}_{\sigma_1^\epsilon,\sigma_2})$. Let us fix an arbitrary strategy σ_2 for Player 2 in the game $\mathcal{G}^{G\vee H}_\rho$. It is enough to show that $[\![J]\!]_\rho(p)-\epsilon\ <E(M^{\langle p,G\vee H\rangle}_{\sigma_1^\epsilon,\sigma_2})$ holds. The Markov chain $M^{\langle p,G\vee H\rangle}_{\sigma_1^\epsilon,\sigma_2}$ in $\mathcal{G}^{G\vee H}_\rho$ can be depicted as in figure 1: from the initial state the edge reaching $\langle p,J\rangle$ is chosen by Player 1, by definition of σ_1^ϵ, and the sub-Markov chain, denoted by M'

and rooted in $\langle p, J\rangle$ which is induced by the strategies σ_1^ϵ and σ_2 from $\langle p, J\rangle$, is reached. Let us define the strategy τ_2, for Player 2 in $\mathcal{G}_\rho^J$ as follows:

$$\tau_2(\vec{s}) \stackrel{\text{def}}{=} \begin{cases} \sigma_2(\langle q, G\vee H\rangle.\vec{t}) & \text{if } \vec{t} \text{ begins in a state } (\langle q, J\rangle) \text{ for } q\in P \\ ? & \text{otherwise} \end{cases}$$

The idea behind the definition of τ_2 is the following: the decision made by τ_2 based on the history $\vec{s}$ having initial state $\langle q, J\rangle$, for $q \in P$ (and in particular having initial state $\langle p, J\rangle$), coincides with the decision made by σ_2 on the history $\langle q, G\vee H\rangle.\vec{s}$, which is indeed a finite path in $\mathcal{G}_\rho^{G\vee H}$. Again the second clause of the above definition, needed to define τ_2 on all game-histories not starting in a state of the form $\langle q, J\rangle$, allows to fix any arbitrary behavior. Note that the strategy τ_1, which in analogy with τ_2 behaves on the history $\vec{s}$ (starting at $\langle q, J\rangle$) as σ_1^ϵ does on the history $\langle q, G\vee H\rangle.\vec{s}$, is τ_1^ϵ, by definition of σ_1^ϵ. By construction we have that the sub-Markov chain M' of $M^{\langle p,G\vee H\rangle}_{\sigma_1^\epsilon,\sigma_2}$, is the Markov chain $M^{\langle p,J\rangle}_{\tau_1^\epsilon,\tau_2}$. By definition of τ_1^ϵ, we have that $E(M') > [\![J]\!]_\rho(p)-\epsilon$, or equivalently by definition of J, $E(M') > [\![G\vee H]\!]_\rho(p)-\epsilon$. To conclude the proof we just need to show that $E(M^{\langle p,G\vee H\rangle}_{\sigma_1^\epsilon,\sigma_2}) = E(M')$. To prove this let us define the bijective function f, from completed paths is M' to completed paths in $M^{\langle p,G\vee H\rangle}_{\sigma_1^\epsilon,\sigma_2}$, defined as $f(\vec{s}) = \langle p, G\vee H\rangle.\vec{s}$. It is easy to check that, by application of lemma 4.15, the equality $\Phi_\rho^J(\vec{s}) = \Phi_\rho^{G\vee H}(f(\vec{s}))$ holds, for all completed paths $\vec{s}\in M'$, where $\Phi_\rho^J(\vec{s})$ and $\Phi_\rho^{G\vee H}$ are the payoff functions of the games $\mathcal{G}_\rho^J$ and $\mathcal{G}_\rho^{G\vee H}$ respectively. Moreover it is routine to check that for every Borel measurable set X of completed paths in M' the following equality holds:

$$\mathcal{M}'(X) = \mathcal{M}^{\langle p,G\vee H\rangle}_{\sigma_1^\epsilon,\sigma_2}(f(X))$$

where $\mathcal{M}'$ is the probability measure over completed paths in $\mathcal{G}_\rho^J$ induced by the Markov chain M', $\mathcal{M}^{\langle p,G\vee H\rangle}_{\sigma_1^\epsilon,\sigma_2}$ is the probability measures over set of completed paths in $\mathcal{G}_\rho^{G\vee H}$ induced by the Markov chain $M^{\langle p,G\vee H\rangle}_{\sigma_1^\epsilon,\sigma_2}$ (see definition 4.7), and the set $f(X)$, which is easily seen to be Borel, is defined as $\{f(\vec{s}) \mid \vec{s}\in X\}$. This concludes our proof for the inequality $[\![F]\!]_\rho(p) \leq Val_\downarrow(\mathcal{G}_\rho^F)(\langle p, F\rangle)$.

Let us now consider the inequality $[\![F]\!]_\rho(p) \geq Val_\uparrow(\mathcal{G}_\rho^F)(\langle p, F\rangle)$, or equivalently by definition 4.11, $[\![F]\!]_\rho(p) \geq \bigsqcap_{\sigma_2}\bigsqcup_{\sigma_1} E(M^{\langle p,F\rangle}_{\sigma_1,\sigma_2})$. We prove this inequality in a way similar to the the previous one by constructing, for each $\epsilon > 0$, a strategy σ_2^ϵ for Player 2 in the game $\mathcal{G}_\rho^{G\vee H}$ which satisfies the inequality $[\![G\vee H]\!]_\rho(p)+\epsilon > \bigsqcup_{\sigma_1} E(M^{\langle p,G\vee H\rangle}_{\sigma_1,\sigma_2^\epsilon})$. By induction hypothesis, we know that the equalities $[\![G]\!]_\rho(p) = Val_\uparrow(\mathcal{G}_\rho^G)(\langle p, G\rangle)$ and $[\![H]\!]_\rho(p) = Val_\uparrow(\mathcal{G}_\rho^H)(\langle p, H\rangle)$ hold. This implies that for every $\epsilon > 0$, Player 2 has a strategy τ_2^ϵ in the game $\mathcal{G}_\rho^G$ and a strategy γ_2^ϵ in the game $\mathcal{G}_\rho^H$ satisfying the following inequalities

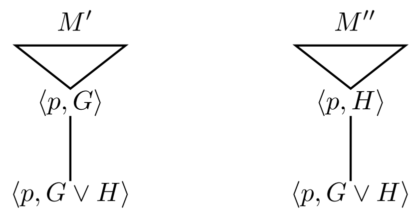

FIGURE 2. Possible shapes of the Markov chain $M^{\langle p, G\vee H\rangle}_{\sigma_1,\sigma_2^\epsilon}$ in $\mathcal{G}^{G\vee H}_\rho$

(1) $[\![G]\!]_\rho + \epsilon > \bigsqcup_{\tau_1} E(M^{\langle p,G\rangle}_{\tau_1,\tau_2^\epsilon})$, and

(2) $[\![H]\!]_\rho + \epsilon > \bigsqcup_{\gamma_1} E(M^{\langle p,H\rangle}_{\gamma_1,\gamma_2^\epsilon})$

respectively. We define the strategy σ_2^ϵ, for Player 2 in the game $\mathcal{G}^{F\vee G}_\rho$ as follows:

$$\sigma_2^\epsilon(\vec{s}) \stackrel{\text{def}}{=} \begin{cases} \tau_2^\epsilon(\langle q, G\rangle.\vec{t}) & \text{if } \vec{s}=\langle q, G\vee H\rangle.\langle q, G\rangle.\vec{t} \ \text{ for some } q\in P \\ \gamma_2^\epsilon(\langle q, H\rangle.\vec{t}) & \text{if } \vec{s}=\langle q, G\vee H\rangle.\langle q, H\rangle.\vec{t} \ \text{ for some } q\in P \\ ? & \text{otherwise} \end{cases}$$

The strategy σ_2^ϵ can be described informally as follows: if the game starts at some initial state of the form $\langle q, G\vee H\rangle$, with $q \in P$ (and in particular on the initial state $\langle p, G\vee H\rangle$), and Player 1 chooses to move to $\langle q, G\rangle$, then the decisions that σ_2 will subsequently take based on a game-history (necessarily) of the form $\langle q, G\vee H\rangle.\langle q, G\rangle.\vec{t}$, will coincide with the decision taken by τ_2^ϵ on the game-history $\langle q, G\rangle.\vec{t}$. Note that $\langle q, G\rangle.\vec{t}$ is indeed a finite path in the game $\mathcal{G}^G_\rho$ (over which τ_2^ϵ is defined). Similarly if the first move of Player 1 on the starting state $\langle q, G\vee H\rangle$ was to take the successor state $\langle q, H\rangle$, then the decisions that σ_2 will subsequently take based on a game-history (necessarily) of the form $\langle q, G\vee H\rangle.\langle q, H\rangle.\vec{t}$, will coincide with the decision taken by γ_2^ϵ on the game-history $\langle q, H\rangle.\vec{t}$. The last clause in the definition of σ_2^ϵ is again just technical, and allows any arbitrary behavior on the game histories $\vec{s}$ not starting on states of the form $\langle q, G\vee H\rangle$, for $q \in P$. We now need to prove that the so defined strategy σ_2^ϵ satisfies the desired inequality $[\![G\vee H]\!]_\rho(p) + \epsilon > \bigsqcup_{\sigma_1} E(M^{\langle p, G\vee H\rangle}_{\sigma_1,\sigma_2^\epsilon})$. Let us fix an arbitrary strategy σ_1 for Player 1 in the game $\mathcal{G}^{G\vee H}_\rho$. We just need to show that $[\![G\vee H]\!]_\rho(p) + \epsilon > E(M^{\langle p, G\vee H\rangle}_{\sigma_1,\sigma_2^\epsilon})$. The Markov chain $M^{\langle p, G\vee H\rangle}_{\sigma_1,\sigma_2^\epsilon}$ in $\mathcal{G}^{G\vee H}_\rho$ can be depicted as one of the two trees of figure 2, depending on whether σ_1 chooses to move to $\langle p, G\rangle$ or $\langle p, H\rangle$ on the first move. By applying the same technique used for proving the first inequality 1, it is possible to show, by the definitions of τ_2^ϵ, γ_2^ϵ and σ_2^ϵ, that $[\![G]\!]_\rho(p) + \epsilon > E(M')$ and similarly $[\![H]\!]_\rho(p) + \epsilon > E(M'')$, and that the expected value $E(M^{\langle p, G\vee H\rangle}_{\sigma_1,\sigma_2^\epsilon}) < ([\![G]\!]_\rho(p)+\epsilon) \sqcup ([\![H]\!]_\rho(p)+\epsilon)$ or equivalently $E(M^{\langle p, G\vee H\rangle}_{\sigma_1,\sigma_2^\epsilon}) < [\![G\vee H]\!]_\rho(p) + \epsilon$, and this concludes the proof.

Case F = ⟨a⟩G. For every process-state $p \in P$ and every interpretation ρ, we have by definition $[\![\langle a\rangle G]\!]_\rho(p) = \bigsqcup\{\overline{[\![G]\!]_\rho}(d) \mid p \xrightarrow{a} d\}$. Let us consider the state $\langle p, \langle a\rangle G\rangle$ of the game $\mathcal{G}_\rho^{\langle a\rangle G}$. This state is in S_1, i.e. under the control of Player 1, which can choose to move to a state (which is under the control of Nature) in the (possibly empty) set $\{\langle d, G\rangle \mid p \xrightarrow{a} d\}$. As a first step, observe that if the set of a-successors of p, i.e. the set $\langle d \mid p \xrightarrow{a} d\}$, is empty, then $[\![\langle a\rangle G]\!]_\rho(p) = 0$, and moreover the game $\mathcal{G}_\rho^{\langle a\rangle G}$ immediately terminates when starting at $\langle p, \langle a\rangle G\rangle$, assigning a reward 0 to Player 1. Therefore in this trivial case, we clearly have $[\![\langle a\rangle G]\!]_\rho(p) = Val_\downarrow(\mathcal{G}_\rho^{\langle a\rangle G})(\langle p, \langle a\rangle G\rangle) = Val_\uparrow(\mathcal{G}_\rho^{\langle a\rangle G})(\langle p, \langle a\rangle G\rangle)$. Let us then assume from now on that $\langle d \mid p \xrightarrow{a} d\}$ is not empty.

We first prove that, $Val_\downarrow(\mathcal{G}_\rho^{\langle a\rangle G})(\langle d, G\rangle) = Val_\uparrow(\mathcal{G}_\rho^{\langle a\rangle G})(\langle d, G\rangle) = \overline{[\![G]\!]_\rho}(d)$, for every process-distribution d such that $p \xrightarrow{a} d$. At the state $\langle d, G\rangle$, Nature chooses to move to the successor state $\langle q, G\rangle$ (with $q \in supp(d)$) with probability $d(q)$; observe that once the state $\langle q, G\rangle$ is reached, the rest of the game is identical to the game $\mathcal{G}_\rho^G$ (starting at $\langle q, G\rangle$). By induction hypothesis, we know that $Val_\downarrow(\mathcal{G}_\rho^G)(\langle q, G\rangle) = Val_\uparrow(\mathcal{G}_\rho^G)(\langle q, G\rangle) = [\![G]\!]_\rho(q)$, for every process state q. Therefore, for every $\epsilon > 0$, Player 1 has a strategy τ_1^ϵ in the game $\mathcal{G}_\rho^G$ such that $\bigsqcap_{\sigma_2} E(M_{\tau_1^\epsilon, \tau_2}^{\langle q, G\rangle}) > [\![G]\!]_\rho(q) - \epsilon$, for every $q \in P$; similarly Player 2 has a strategy τ_2^ϵ in the game $\mathcal{G}_\rho^G$ such that $\bigsqcup_{\tau_1} E(M_{\tau_1, \tau_2^\epsilon}^{\langle q, G\rangle}) < [\![G]\!]_\rho(q) + \epsilon$, for every $q \in P$. We prove that the inequality $Val_\downarrow(\mathcal{G}_\rho^{\langle a\rangle G})(\langle d, G\rangle) \geq \overline{[\![G]\!]_\rho}(d)$ holds, by constructing for every $\epsilon > 0$, a strategy $\sigma_1^{d,\epsilon}$ for Player 1 in the game $\mathcal{G}_\rho^{\langle a\rangle G}$ such that $\bigsqcap_{\sigma_2} E(M_{\sigma_1^{d,\epsilon}, \sigma_2}^{\langle d, G\rangle}) > \overline{[\![G]\!]_\rho}(d) - \epsilon$. The inequality $Val_\uparrow(\mathcal{G}_\rho^{\langle a\rangle G})(\langle d, G\rangle) \leq \overline{[\![G]\!]_\rho}(d)$ can be proven in a similar way. The strategy $\sigma_1^{d,\epsilon}$ is defined as follows:

$$\sigma_1^{d,\epsilon}(\vec{s}) \stackrel{\text{def}}{=} \begin{cases} \tau_1^\epsilon(\vec{t}) & \text{if } \vec{s} = \langle d, G\rangle.\vec{t} \\ ? & \text{otherwise} \end{cases}$$

The strategy $\sigma_1^{d,\epsilon}$ can be described informally as follows: when the game starts at the initial state $\langle d, G\rangle$, after Nature chooses to move to $\langle q, G\rangle$, for some $q \in supp(d)$, the decisions that $\sigma_1^{d,\epsilon}$ will subsequently take based on a game-history (necessarily) of the form $\langle d, G\rangle.\vec{t}$, will coincide with the decision taken by τ_1^ϵ on the game-history $\vec{t}$. Note that $\vec{t}$ is indeed a finite path in the game $\mathcal{G}_\rho^G$ (over which τ_1^ϵ is defined). Therefore $\sigma_1^{d,\epsilon}$ is well defined on all histories starting at $\langle d, G\rangle$. In other words $\sigma_1^{d,\epsilon}$ just waits for the first move of Nature, and then starts behaving as τ_1^ϵ. Again, the question mark in the last clause in the definition of $\sigma_1^{;\epsilon}$ express that on all game histories $\vec{s}$ not starting on states of the form $\langle d, G\rangle$, any arbitrary behavior for $\sigma_1^{d,\epsilon}$ can be fixed. We now need to show that $E(M_{\sigma_1^{d,\epsilon}, \sigma_2}^{\langle d, G\rangle}) > \overline{[\![G]\!]_\rho}(d) - \epsilon$, for every strategy σ_2 for Player 2 in the game $\mathcal{G}_\rho^{\langle a\rangle F}$. The Markov chain $M_{\sigma_1^{d,\epsilon}, \sigma_2}^{\langle d, G\rangle}$ can be depicted as in figure 3, where $\{q_i\}_{i \in I \subseteq \mathbb{N}}$ is the support for d, and the edges

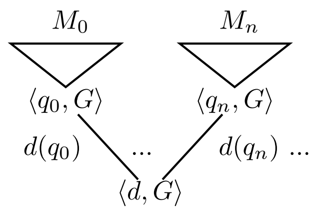

FIGURE 3. Markov chain $M^{\langle d,G\rangle}_{\sigma_1^{d,\epsilon},\sigma_2}$ in $\mathcal{G}_\rho^{\langle a\rangle G}$

leaving the root $\langle d, G\rangle$ have been labeled with their associated probabilities $d(q_i)$. Let us denote by $\mathcal{M}^{\langle d,G\rangle}_{\sigma_1^{d,\epsilon},\sigma_2}$ the probability measure over completed paths induced by $M^{\langle d,G\rangle}_{\sigma_1^{d,\epsilon},\sigma_2}$. By definition 4.7, $\mathcal{M}^{\langle d,G\rangle}_{\sigma_1^{d,\epsilon},\sigma_2}$ assigns probability $d(q_i)$, for $i \in I$, to the clopen set (denoted here by $q_i.\mathcal{P}$) of completed paths having $\langle d, G\rangle.\langle q_i, G\rangle$ as initial prefix. By an argument similar to the one used in the case $F = G \vee H$, it is simple to show that the equality

$$\int_{q_i.\mathcal{P}} \Phi \; d \; \mathcal{M}^{\langle d,G\rangle}_{\sigma_1^{d,\epsilon},\sigma_2} = d(q_i) \cdot \int_{\mathcal{P}} \Phi \; d \; \mathcal{M}_i$$

holds, where Φ is the payoff function of the game $\mathcal{G}_\rho^{\langle a\rangle G}$, and $\mathcal{M}_i$ denotes the probability measure over completed paths induced by the sub-Markov chain M_i, rooted at $\langle q_i, G\rangle$, of $M^{\langle d,G\rangle}_{\sigma_1^{d,\epsilon},\sigma_2}$. Therefore we have that

$$\begin{aligned} E(M^{\langle d,G\rangle}_{\sigma_1^{d,\epsilon},\sigma_2}) &= \int_{\mathcal{P}} \Phi \; d \; \mathcal{M}^{\langle d,G\rangle}_{\sigma_1^{d,\epsilon},\sigma_2} \\ &= \sum_{i \in I} \int_{q_i.\mathcal{P}} \Phi \; d \; \mathcal{M}^{\langle d,G\rangle}_{\sigma_1^{d,\epsilon},\sigma_2} \\ &= \sum_{i \in I} d(q_i) \cdot \big(\int_{\mathcal{P}} \Phi \; d \; \mathcal{M}_i \big) \end{aligned}$$

By an argument similar to the ones used in the previous cases, it is possible to show, exploiting the definition of $\sigma_1^{d,\epsilon}$, that for every $i \in I$, the following inequality holds:

$$E(M_i) = \int_{\mathcal{P}} \Phi \; d \; \mathcal{M}_i > [\![G]\!](q_i) - \epsilon$$

We can conclude the proof by the following sequence of valid inequalities:

$$\begin{aligned} E(M^{\langle d,G\rangle}_{\sigma_1^{d,\epsilon},\sigma_2}) &= \sum_{i\in I}\big(d(q_i)\cdot E(M_i)\big)\\ &> \sum_{i\in I}\big(d(q_i)\cdot(\llbracket G\rrbracket_\rho(q_i)-\epsilon)\big)\\ &= \sum_{i\in I}\big(d(q_i)\cdot(\llbracket G\rrbracket_\rho(q_i)\big)-\epsilon\\ &= \overline{\llbracket G\rrbracket_\rho}(d)-\epsilon \end{aligned}$$

where the second last equality follows from the fact that $\sum_i d(q_i)=1$.

We can now prove the two inequalities $\llbracket\langle a\rangle G\rrbracket(p)_\rho \leq Val_\downarrow(\mathcal{G}^{\langle a\rangle G}_\rho)(\langle p,\langle a\rangle G\rangle)$ and $\llbracket\langle a\rangle G\rrbracket(p)_\rho \geq Val_\uparrow(\mathcal{G}^{\langle a\rangle G}_\rho)(\langle p,\langle a\rangle G\rangle)$.

Let us consider first the inequality $\llbracket\langle a\rangle G\rrbracket(p)_\rho \leq Val_\downarrow(\mathcal{G}^{\langle a\rangle G}_\rho)(\langle p,\langle a\rangle G\rangle)$. We prove this inequality by constructing, for every $\epsilon>0$, a strategy σ_1^ϵ for Player 1 in the game $\mathcal{G}^{\langle a\rangle F}_\rho$ such that the inequality $\llbracket\langle a\rangle G\rrbracket_\rho(p)-\epsilon < \bigsqcap_{\sigma_2} E(M^{\langle p,\langle a\rangle G\rangle}_{\sigma_1^\epsilon,\sigma_2})$ holds. Let d be a process-distribution such that $p\xrightarrow{a}d$, and satisfying $\overline{\llbracket G\rrbracket_\rho}(d) > \llbracket\langle a\rangle G\rrbracket_\rho(p)-\frac{\epsilon}{2}$. The strategy σ_1^ϵ is defined as follows: on the initial state $\langle p,\langle a\rangle G\rangle$ it chooses to move to the state $\langle d,G\rangle$, from which, its subsequent choices made on some history (necessarily) of the form $\langle p,\langle a\rangle G\rangle.\vec{t}$ will coincide with the choices $\tau_1^{d,\frac{\epsilon}{2}}(\vec{t})$, where the strategy $\tau_1^{d,\frac{\epsilon}{2}}$, whose existence was proved before, guarantees Player 1 an expected payoff greater than $\overline{\llbracket G\rrbracket_\rho}(d)-\frac{\epsilon}{2}$. It is easy to check that the so defined strategy σ_1^ϵ, satisfies the desired inequality, since $\overline{\llbracket G\rrbracket_\rho}(d) > \llbracket\langle a\rangle G\rrbracket_\rho(p)-\frac{\epsilon}{2}$ by definition. To conclude the proof we need to show that also the inequality $\llbracket\langle a\rangle G\rrbracket_\rho(p) \geq Val_\uparrow(\mathcal{G}^{\langle a\rangle G}_\rho)(\langle p,\langle a\rangle G\rangle)$ holds. This is done, again, by constructing for every $\epsilon>0$, a strategy σ_2^ϵ for Player 2 in the game $\mathcal{G}^{\langle a\rangle G}_\rho$ satisfying the inequality $\llbracket\langle a\rangle G\rrbracket_\rho(p)+\epsilon > \bigsqcup_{\sigma_1} E(M^{\langle p,\langle a\rangle G\rangle}_{\sigma_1,\sigma_2^\epsilon})$. The strategy σ_2^ϵ is defined as follows: if at the first step of the game starting in $\langle p,\langle a\rangle G\rangle$, Player 1 chooses to move to the state $\langle d,G\rangle$, then Player 2 subsequent decisions based on a history (necessarily) of the form $\langle p,\langle a\rangle G\rangle.\langle d,G\rangle.\vec{t}$, will coincide with $\tau_2^{d,\epsilon}(\langle d,G\rangle.\vec{t})$, where the Player 2 strategy $\tau_2^{d,\epsilon}$, whose existence was proven before, guarantees an expected payoff below $\overline{\llbracket G\rrbracket_\rho}(d)+\epsilon$ to Player 1, from the state $\langle d,G\rangle$. It is routine to check that the so defined strategy σ_2^ϵ satisfies the desired inequality: whatever successor $\langle d,G\rangle$ Player 1 chooses on the initial move, his expected payoff after that move will be bounded above by $\overline{\llbracket G\rrbracket_\rho}(d)+\epsilon$; therefore the expected payoff at the initial state $\langle p,\langle a\rangle G\rangle$ is bounded above by $\bigsqcup\{\overline{\llbracket G\rrbracket_\rho}(d)+\epsilon \mid p\xrightarrow{a}d\}$ or equivalently by $\llbracket\langle a\rangle G\rrbracket(p)+\epsilon$, and this concludes the proof.

Case $\mathbf{F}=\mu\mathbf{X}.\mathbf{G}$. By definition we have that $\llbracket F\rrbracket_\rho(p)=\big(\text{lfp of }\lambda f.\llbracket G\rrbracket_{\rho[f/X]}\big)(p)$, or equivalently, by the Knaster-Tarski theorem, $\llbracket F\rrbracket_\rho(p)=\big(\bigsqcup_\alpha\llbracket G\rrbracket_{\rho^\alpha}\big)(p)$, where ρ^α is defined for each ordinal α as follows:

$$\rho^{\alpha}(Y) = \begin{cases} \rho(Y) & \text{if } X \neq Y \\ \bigsqcup_{\beta<\alpha} \llbracket G \rrbracket_{\rho^{\beta}} & \text{otherwise} \end{cases}$$

Let γ be the smallest ordinal, such that $\llbracket G\rrbracket_{\rho^{\gamma}} = \bigsqcup_{\alpha} \llbracket G\rrbracket_{\rho^{\alpha}} = \rho^{\gamma}(X)$. We need to show that the two inequalities $\llbracket G\rrbracket_{\rho^{\gamma}}(p) \leq Val_{\downarrow}(\mathcal{G}^{\mu X.G}_{\rho})(\langle p, \mu X.G\rangle)$ and $\llbracket G\rrbracket_{\rho^{\gamma}}(p) \geq Val_{\uparrow}(\mathcal{G}^{\mu X.G}_{\rho})(\langle p, \mu X.G\rangle)$ hold. As a first observation, note that $\mathcal{G}^{\mu X.G}_{\rho}$ and $\mathcal{G}^{\mu X.G}_{\rho^{\gamma}}$ are identical, as ρ and ρ^{γ} differs only on the variable X which is not free in $\mu X.G$. Therefore we just need to prove that $\llbracket G\rrbracket_{\rho^{\gamma}}(p) \leq Val_{\downarrow}(\mathcal{G}^{\mu X.G}_{\rho^{\gamma}})(\langle p, \mu X.G\rangle)$ and $\llbracket G\rrbracket_{\rho^{\gamma}}(p) \geq Val_{\uparrow}(\mathcal{G}^{\mu X.G}_{\rho^{\gamma}})(\langle p, \mu X.G\rangle)$ hold. Moreover, since the state $\langle p, \mu X.G\rangle$ has a unique successor $\langle p, G\rangle$ in $\mathcal{G}^{\mu X.G}_{\rho^{\gamma}}$, and since, as stated in lemma 4.15, the payoff function of pLμ games is *prefix independent*, we just need to show that the inequalities $\llbracket G\rrbracket_{\rho^{\gamma}}(p) \leq Val_{\downarrow}(\mathcal{G}^{\mu X.G}_{\rho^{\gamma}})(\langle p, G\rangle)$ and $\llbracket G\rrbracket_{\rho^{\gamma}}(p) \geq Val_{\uparrow}(\mathcal{G}^{\mu X.G}_{\rho^{\gamma}})(\langle p, G\rangle)$ hold. By induction hypothesis we already know that $\llbracket G\rrbracket_{\rho^{\gamma}}(p) = Val_{\downarrow}(\mathcal{G}^{G}_{\rho^{\gamma}})(\langle p, G\rangle) = Val_{\uparrow}(\mathcal{G}^{G}_{\rho^{\gamma}})(\langle p, G\rangle)$, and therefore we just need to prove the following two inequalities:

I) $Val_{\uparrow}(\mathcal{G}^{G}_{\rho^{\gamma}})(\langle p, G\rangle) \geq Val_{\uparrow}(\mathcal{G}^{\mu X.G}_{\rho^{\gamma}})(\langle p, G\rangle)$, and

II) $Val_{\downarrow}(\mathcal{G}^{G}_{\rho^{\gamma}})(\langle p, G\rangle) \leq Val_{\downarrow}(\mathcal{G}^{\mu X.G}_{\rho^{\gamma}})(\langle p, G\rangle)$.

Both inequalities are proven by exploiting the similarities between the two games $\mathcal{G}^{G}_{\rho^{\gamma}}$ and $\mathcal{G}^{\mu X.G}_{\rho^{\gamma}}$ which can be summarized as follows:

- The set $S^{\mu X.G}$ of states of the game $\mathcal{G}^{\mu X.G}_{\rho^{\gamma}}$, is the same as the set S^{G} of states in the game $\mathcal{G}^{G}_{\rho^{\gamma}}$, plus the set of states $\{\langle q, \mu X.G\rangle \mid q \in P\}$. This extra states however play almost no role in the game: $\langle q, \mu X.G\rangle$ is under the control of Player 1, which has a unique forced move to $\langle q, G\rangle$, and after this step, the state $\langle q, \mu X.G\rangle$ can not be visited ever again in any play.
- More significantly, the set of states of the form $\langle q, X\rangle$, for $q \in P$, which are present in both games, are terminal states in $\mathcal{G}^{G}_{\rho^{\gamma}}$ (when the game ends on such a state Player 1 receives a payoff of $\rho^{\gamma}(X)(q)$), and non-terminal states in $\mathcal{G}^{\mu X.G}_{\rho^{\gamma}}$ where Player 1 has only one forced moved to the successor state $\langle q, G\rangle$.

Moreover observe that the priorities assigned to the states of the two games coincide (or at least the can be made to coincide), and similarly the terminal reward assignment of the two games coincide as well, except that in the states of the form $\langle q, X\rangle$ of $\mathcal{G}^{\mu X.G}_{\rho^{\gamma}}$, the terminal reward is not defined since, as observed before, $\langle q, X\rangle$ is not terminal in $\mathcal{G}^{\mu X.G}_{\rho^{\gamma}}$. Given these observations it is clear that any finite path in $\mathcal{G}^{G}_{\rho^{\gamma}}$ is also a finite path in $\mathcal{G}^{\mu X.G}_{\rho^{\gamma}}$. Moreover we define the functions *count* and *tail* from finite paths in $\mathcal{G}^{\mu X.G}_{\rho^{\gamma}}$ to natural numbers and finite paths in $\mathcal{G}^{G}_{\rho^{\gamma}}$ respectively, as follows:

$$count(\vec{s}) = |\{\langle q', X\rangle \in \vec{s} \mid q' \in P\}|$$

and

$$tail(\vec{s}) = \begin{cases} \vec{s'} & \text{if } \vec{s} = \langle q, \mu X.G\rangle.\vec{s} \text{ and } count(\vec{s'}) = 0 \\ \vec{s} & \text{if } count(\vec{s}) = 0 \text{ and } \vec{s} \text{ does not start at } \langle q, \mu X.G\rangle \\ \vec{t} & \text{if } \vec{s} = \vec{s'}.\langle q', X\rangle.\vec{t} \text{ and } count(\vec{t}) = 0 \end{cases}$$

In other words $count(\vec{s})$ gives us the number of occurrences of states of the form $\langle q', X\rangle$ in $\vec{s}$, for $q' \in P$, and the finite path $tail(\vec{s})$ is obtained by removing from $\vec{s}$ by removing the initial state if this is of the form $\langle q, \mu X.G\rangle$ with $q \in P$, and by removing the initial prefix up to the last occurrence of a state of the form $\langle q', X\rangle$ in $\vec{s}$. Note that $tail(\vec{s})$ is indeed a finite path in $\mathcal{G}^{G}_{\rho^\gamma}$.

Let us first prove the inequality I. We do this by constructing, for every $\epsilon > 0$ and for every $k \in \mathbb{N}$, a strategy $\sigma_2^{[k]}$ for Player 2 in the game $\mathcal{G}^{\mu X.G}_{\rho^\gamma}$, satisfying the following inequality: $Val_\uparrow(\mathcal{G}^{G}_{\rho^\gamma})(\langle p, G\rangle) + \frac{\epsilon}{2^k} > \bigsqcup_{\sigma_1} E(M^{\langle p,G\rangle}_{\sigma_1, \sigma_2^{[k]}})$. This strategy is built using the collection of δ-optimal strategies τ_2^δ, with $\delta > 0$, for Player 2 in the game $\mathcal{G}^{G}_{\rho^\gamma}$, i.e. strategies σ_2^δ such that the following inequality holds: $\bigsqcup_{\tau_1} M^{\langle q,G\rangle}_{\tau_1,\tau_2^\delta} < Val_\uparrow(\mathcal{G}^{G}_{\rho^\gamma})(\langle q, G\rangle) + \delta$, for every $q \in P$. The strategy $\sigma_2^{[k]}$ is defined as follows:

$$\sigma_2^{[k]}(\vec{s}) = \begin{cases} \tau_2^{\frac{\epsilon}{2^{k+1}}}(\vec{s}) & \text{if } \vec{s} \text{ starts at } \langle q, G\rangle \text{ for } q \in P, \text{ and } count(\vec{s}) = 0 \\ \sigma_2^{[k+i]}(\vec{t}) & \text{if } count(\vec{s}) = i > 0 \text{ and } \vec{t} = tail(\vec{s}) \\ ? & \text{otherwise} \end{cases}$$

The strategy $\sigma_2^{[k]}$ can be informally described as follows: when the game starts at the initial state $\langle q, G\rangle$, for $q \in P$, $\sigma_2^{[k]}$ initially behaves as the strategy $\tau_2^{\frac{\epsilon}{2^{k+1}}}$; if a state of the form $\langle q', X\rangle$, for $q' \in P$, is ever reached, then Player 2 *improves* his strategy and starts behaving as the strategy σ_2^{k+1} and so on; this means that on a history of the form $\vec{s} = \vec{s'}.\langle q, X\rangle.\vec{t}$, where $\langle q, X\rangle$ is the last occurrence of a state of the form $\langle q', X\rangle$ in $\vec{s}$, $\sigma_2^{[k]}$ decisions on $\vec{s}$, will coincide with $\sigma_2^{[k+i]}$ (or equivalently with $\tau_2^{\frac{\epsilon}{2^{k+i+1}}}$) on $\vec{t}$, where $i = count(\vec{s'}.\langle q', X\rangle)$, i.e. the number of times a state of the form $\langle q', X\rangle$ is visited in $\vec{s}$. In other words Player 2, using the strategy $\sigma_2^{[k]}$, plays in $\mathcal{G}^{\mu X.G}_{\rho}$ as if it was playing in the game $\mathcal{G}^{G}_{\rho^\gamma}$, and every time a state of the form $\langle q, X\rangle$ is reached, he re-starts again (from the unique successor $\langle q, G\rangle$ of $\langle q, X\rangle$) as if he was in $\mathcal{G}^{G}_{\rho^\gamma}$, but with an improved strategy. Again the last clause in the definition of the strategy $\sigma_2^{[k]}$ allows for any arbitrary behavior on all histories not starting in a state of the form $\langle q, G\rangle$, for $q \in P$.

We now prove that for every $k \in \mathbb{N}$, the strategy $\sigma_2^{[k]}$ satisfies the desired inequality $Val_\uparrow(\mathcal{G}^{G}_{\rho^\gamma})(\langle p, G\rangle) + \frac{\epsilon}{2^k} > \bigsqcup_{\sigma_1} E(M^{\langle p,G\rangle}_{\sigma_1, \sigma_2^{[k]}})$. Let us fix an arbitrary strategy σ_1 for Player 1 in the game $\mathcal{G}^{\mu X.G}_{\rho^\gamma}$. We just need to show that the inequality

$Val_{\uparrow}(\mathcal{G}^{G}_{\rho^{\gamma}})(\langle p,G\rangle)+\frac{\epsilon}{2^k} > E(M^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}})$ holds. Let us denote with $\mathcal{X}^n$, for $n\in\mathbb{N}$, and with $\mathcal{X}^\infty$ the sets of completed paths in $M^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}$ that reach a state of the form $\langle q,X\rangle$ exactly n times and infinitely many times respectively. Let $\mathcal{X}^{\leq n}$ be the set $\bigcup_{i\leq n}\mathcal{X}^i$. From the fact that $\Phi(\vec{s})=0$ for every $\vec{s}\in\mathcal{X}^\infty$ (because $\mathbb{P}(\vec{s})$ is odd, where Φ and $\mathbb{P}$ are the payoff function and the priority assignment in the game $\mathcal{G}^{\mu X.G}_{\rho^{\gamma}}$ respectively) we have

$$\begin{aligned}
E(M^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}) &= \int_{\mathcal{P}} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}\\
&= \sum_{n\in\mathbb{N}}\int_{\mathcal{X}^n} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}} + \int_{\mathcal{X}^\infty} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}\\
&= \sum_{n\in\mathbb{N}}\int_{\mathcal{X}^n} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}\\
&= \bigsqcup_{n\in\mathbb{N}}\int_{\mathcal{X}^{\leq n}} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}
\end{aligned}$$

where the second equality follows from countable additivity of the probability measure $\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}$. We now prove by induction on n that for every $k\in\mathbb{N}$ the inequality

$$Val_{\uparrow}(\mathcal{G}^{G}_{\rho^{\gamma}})(\langle p,G\rangle)+\sum_{i\leq n}\frac{\epsilon}{2^{k+i+1}} > \int_{\mathcal{X}^{\leq n}} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}$$

holds. This clearly implies the desired inequality

$$Val_{\uparrow}(\mathcal{G}^{G}_{\rho^{\gamma}})(\langle p,G\rangle)+\tfrac{\epsilon}{2^k} > \bigsqcup_{n\in\mathbb{N}}\int_{\mathcal{X}^{\leq n}} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}$$

since $\displaystyle\bigsqcup_{n\in\mathbb{N}}\sum_{i\leq n}\frac{\epsilon}{2^{k+i+1}}=\frac{\epsilon}{2^k}$.

Suppose by induction hypothesis that the inequality holds for all $n<m$, for some $m\in\mathbb{N}$. The Markov chain $\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}$ can be depicted as in figure 4 where the triangle (denoted by $\mathcal{X}^0$) represents the set of paths in $\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}$ never reaching a state of the form $\langle q,X\rangle$, for $q\in P$, and the finite paths, denoted by $\vec{t_i}$ connecting the root $\langle p,G\rangle$ with the node $\langle q_i,X\rangle$, for $i\in I\subseteq\mathbb{N}$, are the prefixes (up to the first occurrence of a state of the form $\langle q,X\rangle$, namely $\langle q_i,X\rangle$) of all paths $\vec{s}$ in $\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}$ of the form $\vec{s}=\vec{t_i}.\vec{s'}$. The sub-Markov chain rooted after $\vec{t_i}$, i.e. having $\langle q_i,G\rangle$ as initial state, is denoted by M_i.

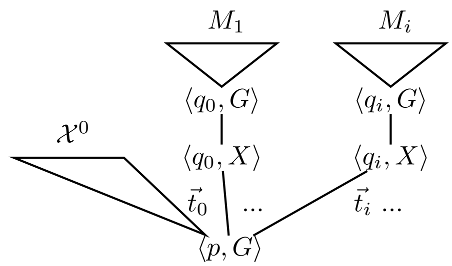

FIGURE 4. Markov chain $M^{\langle p,G\rangle}_{\sigma_1^\epsilon,\sigma_2^{[k]}}$ in $\mathcal{G}^{\mu X.G}_{\rho^\gamma}$

Note that every path $\vec{s}\in\mathcal{X}^{\leq m}$ is either a path in $\mathcal{X}^0$, i.e. does not have any occurrences of states of the form $\langle q, X\rangle$ or is in $\bigcup_{0<j\leq m}\mathcal{X}^j$. Moreover observe that any path $\vec{s}\in\bigcup_{0<j\leq m}\mathcal{X}^j$, i.e. any path that reaches at least once (and at most m times) a state of the form $\langle q, X\rangle$, can be *uniquely* written as the concatenation $\vec{s}=\vec{t_i}.\vec{s'}$ of some finite path $\vec{t_i}$ (ending in the state $\langle q_i, X\rangle$, which is the first occurrence of a state of this shape in $\vec{s}$) and some completed paths $\vec{s'}\in\mathcal{X}^{<m}$, which is necessarily starting in the state $\langle q_i, G\rangle$. Observe that each $\vec{t_i}$, for $i\in I$, is also a finite path in the game $\mathcal{G}^G_{\rho^\gamma}$. Let us denote with $\vec{t_i}.\mathcal{X}^{<m}$, for $i\in I$, the set of paths $\vec{s}\in\bigcup_{0<j\leq m}\mathcal{X}^j$ of the form $\vec{t_i}.\vec{s'}$, with $\vec{s'}\in\mathcal{X}^{<m}$. Given these observations we have that

$$\mathcal{X}^{\leq m}=\mathcal{X}^0\cup\big(\bigcup_{i\in I}\vec{t_i}.\mathcal{X}^{<m}\big)$$

and since the set I is countable, the following equality holds:

$$\int_{\mathcal{X}^{\leq m}}\Phi\ d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}=\int_{\mathcal{X}^0}\Phi\ d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}+\sum_{i\in I}\int_{\vec{t_i}.\mathcal{X}^{<m}}\Phi\ d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}} \tag{1}$$

Moreover, denoting by $\pi(\vec{t_i})$, for $i\in I$, the multiplication of all probabilities appearing in the probabilistic steps of the path $\vec{t_i}$, it is simple to check that the following equality holds:

$$\int_{\vec{t_i}.\mathcal{X}^{<m}}\Phi\ d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}=\pi(\vec{t_i})\cdot\int_{\mathcal{X}^{<m}}\Phi\ d\mathcal{M}_i \tag{2}$$

where, as usual, $\mathcal{M}_i$ denotes the probability measure over completed paths induced by the sub-Markov chain M_i. Therefore, from 1 and 2, we have that the following equality holds:

$$\int_{\mathcal{X}^{\leq m}}\Phi\ d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}=\int_{\mathcal{X}^0}\Phi\ d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}+\sum_{i\in I}\pi(\vec{t_i})\cdot\big(\int_{\mathcal{X}^{<m}}\Phi\ d\mathcal{M}_i\big) \tag{3}$$

Let us define, for every every $i \in I$, the strategies σ_1^i and σ_2^i for Player 1 and Player 2 respectively in the game $\mathcal{G}_{\rho^\gamma}^{\mu X.G}$ as follows:

$$\sigma_1^i(\vec{s}) = \begin{cases} \sigma_1(\vec{t_i}.\vec{s}) & \text{if } \vec{t_i}.\vec{s} \text{ is a finite path in } \mathcal{G}_{\rho^\gamma}^{\mu X.G}, \text{ i.e. if } \vec{s} \text{ extends } \vec{t_i} \\ ? & \text{otherwise} \end{cases}$$
$$\sigma_2^i(\vec{s}) = \begin{cases} \sigma_2^{[k]}(\vec{t_i}.\vec{s}) & \text{if } \vec{t_i}.\vec{s} \text{ is a finite path in } \mathcal{G}_{\rho^\gamma}^{\mu X.G}, \text{ i.e. if } \vec{s} \text{ extends } \vec{t_i} \\ ? & \text{otherwise} \end{cases}$$

In other words σ_1^i decisions on the history $\vec{s}$ coincide with σ_1's decisions on $\vec{t_i}.\vec{s}$ if $\vec{s}$ is a possible extension of $\vec{t_i}$; otherwise any arbitrary behavior is allowed. Similarly σ_2^i decisions on the history $\vec{s}$ coincide with $\sigma_2^{[k]}$'s decisions on $\vec{t_i}.\vec{s}$ if $\vec{s}$ is a possible extension of $\vec{t_i}$; otherwise any arbitrary behavior is allowed. From these definitions, it follows that the sub-Markov chain M_i of $M^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}$ rooted at $\vec{t_i}$, is the Markov chain $M^{\langle q_i,G\rangle}_{\sigma_1^i,\sigma_2^i}$. Moreover it follows, from definition of $\sigma_2^{[k]}$, that for every $i \in I$ and for every possible extension $\vec{s}$ of $\vec{t_i}$, that the equalities $\sigma_2^i(\vec{s}) = \sigma_2^{[k+1]}(\vec{s}) = \sigma_2^k(\vec{t_i}.\vec{s})$ hold. Therefore we can equate the following expressions:

$$\int_{\mathcal{X}^{<m}} \Phi \, d\mathcal{M}_i = \int_{\mathcal{X}^{<m}} \Phi \, d\mathcal{M}^{\langle q_i,G\rangle}_{\sigma_1^i,\sigma_2^{[k+1]}} \tag{4}$$

By induction hypothesis we know that

$$\int_{\mathcal{X}^{<m}} \Phi \, d\mathcal{M}^{\langle q_i,G\rangle}_{\sigma_1^i,\sigma_2^{[k+1]}} < Val_\uparrow(\mathcal{G}^G_{\rho^\gamma})(\langle q_i, G\rangle) + \sum_{j<m} \frac{\epsilon}{2^{(k+1)+j+1}} \tag{5}$$

and so, from 4 and 5, the following inequality holds:

$$\int_{\mathcal{X}^{<m}} \Phi \, d\mathcal{M}_i < \big(Val_\uparrow(\mathcal{G}^G_{\rho^\gamma})(\langle q_i, G\rangle) + \sum_{j<m} \frac{\epsilon}{2^{(k+1)+j+1}}\big) \tag{6}$$

Let us now define the two strategies τ_1 and τ_2, for Player 1 and Player 2 respectively, in the game $\mathcal{G}^G_{\rho^\gamma}$ as $\tau_1(\vec{s}) = \sigma_1(\vec{s})$ and $\tau_2(\vec{s}) = \sigma_2^{[k]}(\vec{s})$. Note that this is a good definition since, as observed above, any finite path $\vec{s}$ in $\mathcal{G}^G_{\rho^\gamma}$ is also a finite path in $\mathcal{G}^{\mu X.G}_{\rho^\gamma}$, over which σ_1 and $\sigma_2^{[k]}$ are defined. The intuition about the strategies τ_1 and τ_2 is that they behave as the strategies σ_1 and $\sigma_2^{[k]}$; if a state of the form $\langle q, X\rangle$ is ever reached, the game $\mathcal{G}^G_{\rho^\gamma}$ terminates (as $\langle q, X\rangle$ is terminal in $\mathcal{G}^G_{\rho^\gamma}$) and therefore no further action has to be taken. Therefore the Markov chain $M^{\langle p,G\rangle}_{\tau_1,\tau_2}$, in $\mathcal{G}^G_{\rho^\gamma}$, which can be depicted as in figure 5, is identical to the Markov chain $M^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}$ in $\mathcal{G}^{\mu X.G}_{\rho^\gamma}$ except that each node of the form $\langle q_i, X\rangle$, for $i \in I$, is a leaf in $M^{\langle p,G\rangle}_{\tau_1,\tau_2}$, whereas in $M^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}$ it is the root of the sub-Markov chain M_i. Recall

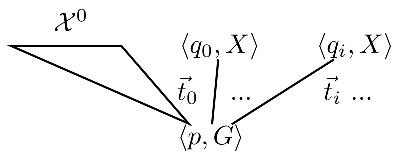

FIGURE 5. Markov chain $M^{\langle p,G\rangle}_{\tau_1,\tau_2}$ in $\mathcal{G}^G_{\rho^\gamma}$

that whenever a play in the game $\mathcal{G}^G_{\rho^\gamma}$ ends in one of the terminal states $\langle q_i, X\rangle$, Player 1 receives a payoff of $\rho^\gamma(X)(q_i)$. Given these observations, the following equality holds:

$$E(M^{\langle p,G\rangle}_{\tau_1,\tau_2}) = \int_{\mathcal{P}} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\tau_1,\tau_2} = \int_{\mathcal{X}^0} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\tau_1,\tau_2} + \sum_{i\in I} \pi(\vec{t}_i) \cdot \big(\rho^\gamma(X)(q_i)\big) \tag{7}$$

As observed above, since $M^{\langle p,G\rangle}_{\tau_1,\tau_2}$ and $\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}}$ are identical up to the first occurrences of states states of the form $\langle q_i, X\rangle$, for $i \in I$, we have that the following equality holds:

$$\int_{\mathcal{X}^0} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\tau_1,\tau_2} = \int_{\mathcal{X}^0} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}} \tag{8}$$

beacuse the set $\mathcal{X}^0$ is the set of path never reaching a state of the form $\langle q_i, X\rangle$.

Moreover, since $\rho^\gamma(X)(q_i) = [\![G]\!]_{\rho^\gamma}(q_i)$ by definition of γ, since by inductive hypothesis (on the structure of pLμ formuale) we know that $[\![G]\!]_{\rho^\gamma}(q_i) = Val_\uparrow(\mathcal{G}^G_{\rho^\gamma})(\langle q_i, G\rangle)$, and since that inequality 6 holds, we have that

$$\int_{\mathcal{X}^{<m}} \Phi \; d\mathcal{M}_i < \big(\rho^\gamma(X)(q_i) + \sum_{j<m} \frac{\epsilon}{2^{(k+1)+j+1}}\big) \tag{9}$$

which in turn implies the following inequality:

$$\sum_{i\in I} \pi(\vec{t}_i) \cdot \int_{\mathcal{X}^{<m}} \Phi \; d\mathcal{M}_i < \sum_{i\in I} \pi(\vec{t}_i) \cdot \big(\rho^\gamma(X)(q_i) + \sum_{j<m} \frac{\epsilon}{2^{(k+1)+j+1}}\big) \tag{10}$$

from which

$$\sum_{i\in I} \pi(\vec{t}_i) \cdot \int_{\mathcal{X}^{<m}} \Phi \; d\mathcal{M}_i < \big(\sum_{i\in I} \pi(\vec{t}_i) \cdot \big(\rho^\gamma(X)(q_i)\big) + \sum_{j<m} \frac{\epsilon}{2^{(k+1)+j+1}} \tag{11}$$

follows because $\sum_{i\in I} \pi(\vec{t}_i) \leq 1$. From 3, 8 and 11 we can derive the following inequality:

$$\int_{\mathcal{X}^{\leq m}} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}} < \big(\int_{\mathcal{P}} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\tau_1,\tau_2}\big) + \sum_{j<m} \frac{\epsilon}{2^{(k+1)+j+1}}. \tag{12}$$

To conclude the proof we need the following last observation: the strategy τ_2 (which behaves as $\sigma_2^{[k]}$ up to the first occurrence of a state of the form $\langle q, X\rangle$) coincides by definition of $\sigma_2^{[k]}$ with the strategy $\tau_2^{\frac{\epsilon}{2^{k+1}}}$ which satisfies the inequality $\bigsqcup_{\tau_1} E(M^{\langle p,G\rangle}_{\tau_1,\tau_2^{\frac{\epsilon}{2^{k+1}}}}) < Val_{\uparrow}(\mathcal{G}^G_{\rho^\gamma})(\langle p, G\rangle) + \frac{\epsilon}{2^{k+1}}$. Therefore from 12, we have that the inequality

$$\int_{\mathcal{X}^{\leq m}} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}} < \big(Val_{\uparrow}(\mathcal{G}^G_{\rho^\gamma})(\langle p, G\rangle) + \frac{\epsilon}{2^{k+1}}\big) + \sum_{j<m} \frac{\epsilon}{2^{(k+1)+j+1}}.$$

or equivalently the desired inequality

$$\int_{\mathcal{X}^{\leq m}} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1,\sigma_2^{[k]}} < Val_{\uparrow}(\mathcal{G}^G_{\rho^\gamma})(\langle p, G\rangle) + \sum_{j\leq m} \frac{\epsilon}{2^{k+j+1}}.$$

holds, and this concludes the proof of the inequality I.

Let us now prove the inequality II: $Val_{\downarrow}(\mathcal{G}^G_{\rho^\gamma})(\langle p, G\rangle) \leq Val_{\downarrow}(\mathcal{G}^{\mu X.G}_{\rho^\gamma})(\langle p, G\rangle)$. For proving this inequality we use a different technique, namely we prove that for every ordinal α, the inequality

$$Val_{\downarrow}(\mathcal{G}^G_{\rho^\alpha})(\langle p, G\rangle) \leq Val_{\downarrow}(\mathcal{G}^{\mu X.G}_{\rho^\gamma})(\langle p, G\rangle)$$

holds. This clearly implies the desired inequality taking $\alpha = \gamma$. We prove this, by transfinite induction on the ordinals, by constructing for every $\epsilon > 0$ a strategy $\sigma_1^{\alpha,\epsilon}$ for Player 1 in the game $\mathcal{G}^{\mu X.G}_{\rho^\gamma}$ such that the inequality

$$Val_{\downarrow}(\mathcal{G}^G_{\rho^\alpha})(\langle p, G\rangle) - \epsilon < \textstyle\bigsqcap_{\sigma_2} E(M^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2})$$

holds. Suppose the property holds for every ordinal $\beta < \alpha$. This implies that for every ordinal $\beta < \alpha$ and for every $\delta > 0$, there exists a strategy $\sigma_1^{\beta,\delta}$ for Player 1 in the game $\mathcal{G}^{\mu X.G}_{\rho^\gamma}$ such that the following inequality holds for every $q \in P$:

$$Val_{\downarrow}(\mathcal{G}^G_{\rho^\beta})(\langle q, G\rangle) + \delta < \textstyle\bigsqcap_{\sigma_2} E(M^{\langle q,G\rangle}_{\sigma_1^{\beta,\delta},\sigma_2})$$

Let $\tau^{\frac{\epsilon}{2}}$ be a $\frac{\epsilon}{2}$-optimal strategy for Player 1 in the game $\mathcal{G}^G_{\rho^\alpha}$, i.e. such that the inequality

$$Val_{\downarrow}(\mathcal{G}^G_{\rho^\alpha})(\langle p, G\rangle) - \frac{\epsilon}{2} < \textstyle\bigsqcap_{\tau_2} E(M^{\langle p,G\rangle}_{\tau_1^{\frac{\epsilon}{2}},\tau_2})$$

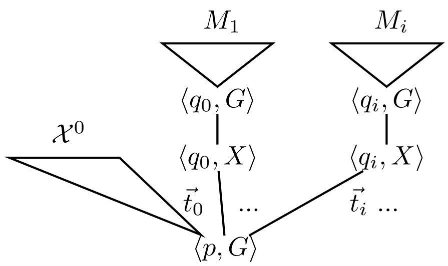

FIGURE 6. Markov chain $M^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2}$ in $\mathcal{G}^{\mu X.G}_{\rho^\gamma}$

is satisfied.

Let us now consider the case where α is a successor ordinal, i.e. $\alpha = \beta + 1$. The strategy $\sigma_1^{\alpha,\epsilon}$ is defined in this case as follows:

$$\sigma_1^{\alpha,\epsilon}(\vec{s}) = \begin{cases} \tau_1^{\frac{\epsilon}{2}}(\vec{s}) & \text{if } \vec{s} \text{ starts at } \langle p, G\rangle \text{ and } count(\vec{s}) = 0 \\ \sigma_1^{\beta,\frac{\epsilon}{2}}(\vec{t}) & \text{if } \vec{s} = \vec{s'}.\langle q, X\rangle.\vec{t} \text{ and } count(\vec{s}) = count(\vec{t}) + 1 \\ ? & \text{otherwise} \end{cases}$$

The strategy $\sigma_2^{\alpha,\epsilon}$ can be informally described as follows: when the game starts at the initial state $\langle p, G\rangle$, $\sigma_2^{\alpha,\epsilon}$ initially behaves as the strategy $\tau_1^{\frac{\epsilon}{2}}$; if a state of the form $\langle q, X\rangle$, for $q \in P$, is ever reached, then Player 1 will play the rest of the game following the strategy $\sigma_1^{\beta,\frac{\epsilon}{2}}$; this means that on a history of the form $\vec{s} = \vec{s'}.\langle q, X\rangle.\vec{t}$, where $\langle q, X\rangle$ is the first occurrence of a state of the form $\langle q', X\rangle$ in $\vec{s}$, $\sigma_1^{\alpha,\epsilon}$ on $\vec{s}$ coincides with $\sigma_1^{\beta,\frac{\epsilon}{2}}$ on $\vec{t}$. Again the last clause in the definition of the strategy $\sigma_2^{\alpha,\epsilon}$ allows for any arbitrary behavior on all histories not considered in the first two cases.

We now need to prove that the strategy $\sigma_2^{\alpha,\epsilon}$ satisfies the desired inequality. This is done using a similar methodology to the one adopted for proving the inequality I. Let us fix an arbitrary strategy σ_2 for Player 2 in the game $\mathcal{G}^{\mu X.G}_{\rho^\gamma}$. We need to show that the inequality

$$Val_{\downarrow}(\mathcal{G}^{G}_{\rho^\alpha})(\langle p, G\rangle) - \epsilon < E(M^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2})$$

holds. The Markov chain $M^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2}$ can be depicted, adopting the same notation used in figure 4, as in figure 6. By consideration analogous to the ones adopted in the proof on inequality I, we have that the following equality holds:

$$E(M^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2}) = \int_{\mathcal{X}^0} \Phi \, d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2} + \sum_{i \in I} \int_{O_{\vec{t}_i}} \Phi \, d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2} \tag{13}$$

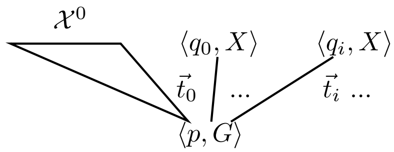

FIGURE 7. Markov chain $M^{\langle p,G\rangle}_{\tau_1,\tau_2}$ in $\mathcal{G}^G_{\rho^\alpha}$

where $O_{\vec{t}_i}$ denotes the set of completed paths in $M^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2}$ having $\vec{t_i}$ as initial prefix, and Φ is the payoff function of the game $\mathcal{G}^{\mu X.G}_{\rho^\gamma}$. Again, it is simple to check the following equality

$$\int_{O_{\vec{t}_i}} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2} = \pi(\vec{t_i}) \cdot E(M_i) \tag{14}$$

and by exploiting the definition of $\sigma_1^{\alpha,\epsilon}$, which once on the history $\vec{t_i}$ starts behaving as the strategy $\sigma_1^{\beta,\frac{\epsilon}{2}}$, it is possible to show that the inequality

$$E(M_i) > Val_{\downarrow}(\mathcal{G}^G_{\rho^\beta})(\langle q_i, G\rangle) - \frac{\epsilon}{2} \tag{15}$$

holds. By induction hypothesis (on the structure of the formula), we know that $Val_{\downarrow}(\mathcal{G}^G_{\rho^\beta})(\langle q_i, G\rangle) = [\![G]\!]_{\rho^\beta}(q_i)$, which implies, by definition of ρ^α, that the inequality

$$E(M_i) > \rho^\alpha(X)(q_i) - \frac{\epsilon}{2} \tag{16}$$

holds. Therefore, by 13, 14 and 16, we have that the following inequality holds:

$$E(M^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2}) > \int_{\mathcal{X}^0} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2} + \big(\sum_{i\in I} \pi(\vec{t_i}) \cdot \rho^\alpha(X)(q_i)\big) - \frac{\epsilon}{2} \tag{17}$$

Let us now define the two strategies τ_1 and τ_2, for Player 1 and Player 2 respectively, in the game $\mathcal{G}^G_{\rho^\alpha}$ as $\tau_1(\vec{s})=\sigma_1^{\alpha,\epsilon}(\vec{s})$ and $\tau_2(\vec{s})=\sigma_2(\vec{s})$. Note that this is a good definition since any finite path $\vec{s}$ in $\mathcal{G}^G_{\rho^\alpha}$ is also a finite path in $\mathcal{G}^{\mu X.G}_{\rho^\gamma}$, over which $\sigma_1^{\alpha,\epsilon}$ and σ_2 are defined. The intuition about the strategies τ_1 and τ_2 is, as usual, that they behave as the strategies $\sigma_1^{\alpha,\epsilon}$ and σ_2; if a state of the form $\langle q, X\rangle$ is ever reached, the game $\mathcal{G}^G_{\rho^\alpha}$ terminates (as $\langle q, X\rangle$ is terminal in $\mathcal{G}^G_{\rho^\alpha}$) and therefore no further action has to be taken. Therefore the Markov chain $M^{\langle p,G\rangle}_{\tau_1,\tau_2}$, in $\mathcal{G}^G_{\rho^\alpha}$, which can be depicted as in figure 7, is identical to the Markov chain $M^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2}$ in $\mathcal{G}^{\mu X.G}_{\rho^\gamma}$ except that each node of the form $\langle q_i, X\rangle$, for $i\in I$, is a leaf in $M^{\langle p,G\rangle}_{\tau_1,\tau_2}$, whereas in $M^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2}$ it is the root of the sub-Markov chain M_i. Recall that whenever a play in the game $\mathcal{G}^G_{\rho^\alpha}$ ends in one of the terminal states $\langle q_i, X\rangle$, Player 1 receives a payoff of $\rho^\alpha(X)(q_i)$. Given these observations, the following equality holds:

$$E(M^{\langle p,G\rangle}_{\tau_1,\tau_2}) = \int_{\mathcal{X}^0} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\tau_1,\tau_2} + \sum_{i\in I} \pi(\vec{t}_i)\cdot \big(\rho^\alpha(X)(q_i)\big) \tag{18}$$

As observed above, since $M^{\langle p,G\rangle}_{\tau_1,\tau_2}$ and $\mathcal{M}^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2}$ are identical up to the first occurrences of states states of the form $\langle q_i, X\rangle$, for $i \in I$, we have that the following equality holds:

$$\int_{\mathcal{X}^0} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\tau_1,\tau_2} = \int_{\mathcal{X}^0} \Phi \; d\mathcal{M}^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2} \tag{19}$$

beacuse the set $\mathcal{X}^0$ is the set of path never reaching a state of the form $\langle q_i, X\rangle$. Given the inequalities 17, 18 and 19 the following inequality follows:

$$E(M^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2}) > E(M^{\langle p,G\rangle}_{\tau_1,\tau_2}) - \frac{\epsilon}{2} \tag{20}$$

To conclude the proof we need, again, the following last observation: the strategy τ_1 (which behaves as $\sigma_2^{\alpha,\epsilon}$ up to the first occurrence of a state of the form $\langle q, X\rangle$) coincides by definition of $\sigma_1^{\alpha,\epsilon}$ with the strategy $\tau_1^{\frac{\epsilon}{2}}$ which satisfies the inequality $\bigsqcap_{\tau_2} E(M^{\langle p,G\rangle}_{\tau_1^{\frac{\epsilon}{2}},\tau_2}) > Val_{\downarrow}(\mathcal{G}^G_{\rho^\alpha})(\langle p,G\rangle) - \frac{\epsilon}{2}$. Therefore from inequality 20, we have that the inequality

$$E(M^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2}) > Val_{\uparrow}(\mathcal{G}^G_{\rho^\alpha})(\langle p,G\rangle) - \frac{\epsilon}{2} - \frac{\epsilon}{2} \tag{21}$$

holds, and this concludes the proof for the case $\alpha = \beta + 1$.

Let us now consider the case when α is a limit ordinal $\alpha = \bigcup_{\beta<\alpha}\beta$. Let β' the least ordinal $\beta' < \alpha$, such that the inequality

$$[\![G]\!]_{\rho^{\beta'}}(q) > \rho^\alpha(X)(q) - \tfrac{\epsilon}{4}$$

holds for every $q \in P$. This ordinal exists, since by definition of ρ^α, $\rho^\alpha(X)(q) = \bigsqcup_{\beta<\alpha}[\![G]\!]_{\rho^{\beta'}}(q)$. The strategy $\sigma_1^{\alpha,\epsilon}$ is defined in this case as follows:

$$\sigma_1^{\alpha,\epsilon}(\vec{s}) = \begin{cases} \tau_1^{\frac{\epsilon}{2}}(\vec{s}) & \text{if } \vec{s} \text{ starts at } \langle p,G\rangle \text{ and } count(\vec{s})=0 \\ \sigma_1^{\beta',\frac{\epsilon}{4}}(\vec{t}) & \text{if } \vec{s}=\vec{s'}.\langle q,X\rangle.\vec{t} \text{ and } count(\vec{s})=count(\vec{t})+1 \\ ? & \text{otherwise} \end{cases}$$

To prove that $\sigma_1^{\alpha,\epsilon}$ satisfies the desired inequality

$$Val_{\downarrow}(\mathcal{G}^G_{\rho^\alpha})(\langle p,G\rangle) - \epsilon < \bigsqcap_{\sigma_2} E(M^{\langle p,G\rangle}_{\sigma_1^{\alpha,\epsilon},\sigma_2})$$

when can apply, without any change, the same technique adopted for the case α successor ordinal. This concludes the proof for the case $F = \mu X.G$.

$\square$

Acknowledgments. The author would like to thank Alex Simpson for helpful discussions.

References

[1] M. Dam. CTL* and ECTL* as fragments of the modal μ-calculus. In *Theoretical Computer Science, Volume 126, Issue 1*, pages 77,96. Elsevier Science B.V., 1994.

[2] 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.

[3] E. A. Emerson and C. S. Jutla. Tree automata, mu-calculus and determinacy. In *SFCS '91: Proceedings of the 32nd annual symposium on Foundations of computer science*, pages 368–377. IEEE Computer Society, 1991.

[4] D. Fischer, E. Grädel, and L. Kaiser. Model checking games for the quantitative μ-calculus. In *Theory of Computing Systems*. Springer New York, 2009.

[5] M. Huth and M. Kwiatkowska. Quantitative analysis and model checking. In *LICS 1997*, page 111, Washington, DC, USA, 1997. IEEE Computer Society.

[6] D. Janin and I. Walukiewicz. On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic. *Lecture Notes in Computer Science*, 1119:263–277, 1996.

[7] D. Kozen. Results on the propositional mu-calculus. In *Theoretical Computer Science*, pages 333–354, 1983.

[8] A. Maitra and W. Sudderth. Borel stochastic games with lim sup payoff. *Annals of Probability*, 21(2):861–885, 1993.

[9] D. A. Martin. The determinacy of Blackwell games. In *Journal of Symbolic Logic Volume 63, Issue 4, 1565-1581*, 1998.

[10] A. McIver and C. Morgan. Results on the quantitative μ-calculus qMμ. *ACM Trans. Comput. Logic*, 8(1):3, 2007.

[11] 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.

[12] R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, Laboratory for Computer Science, Massachusetts Institute of Technology, 1995.

[13] C. Stirling. *Modal and temporal logics for processes*. Springer (Texts in Computer Science), 2001.

[14] W. Zielonka. Perfect-information stochastic parity games. In *FoSSaCS*, volume 2987, pages 499–513, 2004.

Communicated by (The editor will be set by the publisher).
(The dates will be set by the publisher).