

Games, probability and the quantitative μ -calculus $qM\mu$

AK McIver¹ and CC Morgan²

¹ Dept. Computer Science, Macquarie University, NSW 2109 Australia;
anabel@ics.mq.edu.au

² Dept. Comp. Sci. & Eng., University of New South Wales, NSW 2052 Australia;
carrollm@cse.unsw.edu.au

Abstract. The μ -calculus is a powerful tool for specifying and verifying transition systems, including those with demonic (universal) and angelic (existential) choice; its *quantitative* generalisation $qM\mu$ extends that to *probabilistic* choice. We show for a finite-state system that the straightforward *denotational* interpretation of the quantitative μ -calculus is equivalent to an *operational* interpretation given as a turn-based gambling game between two players.

Kozen defined the standard Boolean-typed calculus denotationally; later Stirling gave it an operational interpretation as a turn-based game between two players, and showed the two interpretations equivalent. By doing the same for the quantitative real-typed calculus, we set it on a par with the standard calculus, in that it too can benefit from a solid interface linking the logical and operational frameworks.

Stirling's game analogy, as an aid to intuition, continues in the more general context to provide a surprisingly practical specification tool, meeting for example Vardi's challenge to "figure out the meaning of $AF AX p$ " as a branching-time formula.

We also show that *memoriless strategies suffice* for achieving the *minimax* value of a quantitative game, when the state space is finite.

1 Introduction

The standard μ -calculus, introduced by Kozen [15], extends Boolean dynamic program logic by the introduction of least (μ) and greatest (ν) fixed point operators. Its proof system is applicable to both infinite and finite state spaces; recent results [33] have established a complete axiomatisation; and it can be specialised to temporal logic.

General μ -calculus expressions can be difficult to use, however, because in all but the simplest cases they are not easy on the intuition, especially for example with the nesting of alternating fixed points: even the more specialised temporal properties (particularly "branching-time properties") are notoriously difficult to specify [31]. Stirling's "two-player-game" interpretation alleviates this problem by providing an alternative and operational view [30].

The *quantitative* modal μ -calculus acts over *probabilistic* transition systems, extending the above from Boolean- to real values; and it would benefit just as

much from having two complementary interpretations. We provide them, and show that over a finite state space they are equivalent: one interpretation (defined earlier [19, 12, 7]) generalises Kozen’s; the other (defined here) generalises Stirling’s.

The contribution of this paper is the definition of the Stirling-style quantitative interpretation, and the proof of its equivalence to the Kozen-style quantitative interpretation. We show also that memoriless strategies suffice, for both interpretations, again when the state space is finite.

The Kozen-style quantitative interpretation is based on the extension [14, 24] of Dijkstra/Hoare logic to probabilistic/demonic programs (corresponding to the \forall modality): it is a real-valued logic based on greatest pre-expectations of random variables, rather than weakest preconditions of predicates. It can express “the probability of achieving a postcondition” (since the probability of an event is the expected value of its characteristic function), but it applies more generally to other cost-based properties besides. Although the former may be more intuitive, the extra generality of a “logic of expectations” seems necessary for compositionality [17].

Converting predicates “wholesale” from Boolean- to real-valued state functions — due originally to Kozen [13] and extended by us to include demonic (universal) [24] and angelic (existential) [16] nondeterminism — contrasts with probabilistic logics using “threshold functions” [2, 25] that mix Boolean and numeric arguments: the uniformity in our case means that standard Boolean identities in branching-time temporal logic [1] suggest corresponding quantitative laws for us [20], and so we get a powerful collection of algebraic properties “for free”. The logical “implies” relation between Booleans is replaced by the standard “ \leq ” order on the reals; **false** and **true** become 0 and 1; and fixed points are then associated with monotonic *real*-valued functions. The resulting arithmetic logic is applicable to a restricted class of real-valued functions, and we recall its definition in Sec. 3.

Our Stirling-style quantitative interpretation is operational, and is based on his earlier strategy-based game metaphor for the standard μ -calculus. In our richer context, however, we must distinguish nondeterministic choice from probabilistic choice: the former continues to be represented by the two players’ strategies; but the latter is represented by the new feature that we make the players *gamble*. In Sec. 4 we set out the details.

In Sec. 5 we give an example of the full use of the quantitative aspects of the calculus, beyond simply calculating probabilities.

The main mathematical result of this paper is given in Sec. 6. Stirling showed that for non-probabilistic formulae the Boolean value of the Kozen interpretation corresponds to the existence of a winning strategy in his game interpretation. In our case, strategies in the game must become “optimal” rather than “winning”; and the correspondence is now between a formula’s value (since it denotes a real

number, in the Kozen interpretation) and the expected winnings from the zero-sum gambling game (of the Stirling interpretation). Since the gambling game described by a formula is a “minimax”, we must show it to be well-defined (equal to the “maximin”): in fact we show that both the minimax and the maximin of the game are equal to the Kozen-style denotation of the formula that generated it. We also prove that memoriless strategies suffice.

The benefit of this approach is to set the quantitative μ -calculus on a par with standard μ -calculus in that a suitable form of “logical validity” corresponds exactly to an operational interpretation. As with standard μ -calculus, a specifier can build his intuitions into a game, and can then use the general features of the logic to prove properties about the specific application. For example, the *sublinearity* [24] of $qM\mu$ — the quantitative generalisation of the *conjunctivity* of standard modal algebras — has been used in its quantitative temporal subset qTL to prove a number of algebraic laws strongly reminiscent of those holding in standard branching-time temporal logic [20].

Preliminary experiments have shown that the proof system is very effective for unravelling the intricacies of distributed protocols [28, 21]. Moreover it provides an attractive proof framework for Markov decision processes [23, 9] and indeed many of the problems there have a succinct specification as μ -calculus formulae, as the example of Sec. 5 illustrates. In “reachability-style problems” [6], proof-theoretic methods based on the logic presented here have produced very direct arguments related to the abstraction of probabilities [22], and even more telling is that the logic is applicable even in infinite state spaces [6]. All of which is to suggest that further exploration of $qM\mu$ will continue to be fruitful.

In the following we shall assume generally that S is a countable state space (though for the principal result we restrict to finiteness, in Sec. 6). If f is a function with domain X then by $f.x$ we mean f applied to x , and $f.x.y$ is $(f.x).y$ where appropriate. We denote the set of discrete probability sub-distributions over S by \overline{S} : it is the set of functions from S into the real interval $[0, 1]$ that sum to no more than 1; and if A is a random variable with respect to some probability space, and δ is some probability sub-distribution, we write $\int_{\delta} A$ for the expected value of A with respect to δ .³ In the special case that δ is in \overline{S} and A is a (bounded) real-valued function on S , in fact $\int_{\delta} A$ is equal to $\sum_{s \in S} A.s \times \delta.s$.

2 Probabilistic transition systems and μ -calculus

In this section we set out the language, together with some details about the probabilistic systems over which the formulae are to be interpreted.

³ Normal mathematical practice is to write $\int A d\delta$, but that greatly confuses the roles of bound and free variables.

Formulae in the language (in positive⁴ form) are constructed as follows, where A , K and G all represent constant terms in the uninterpreted formulae:

$$\phi \hat{=} X \mid A \mid \langle K \rangle \phi \mid [K] \phi \mid \phi_1 \sqcap \phi_2 \mid \phi_1 \sqcup \phi_2 \mid \phi_1 \triangleleft G \triangleright \phi_2 \mid (\mu X \cdot \phi) \mid (\nu X \cdot \phi) .$$

It is well-known that these formulae can be used to express complex path-properties of computational sequences; in this paper we shall interpret the formulae over sequences based on generalised probabilistic transitions⁵ in what we call $\mathcal{R}.S$, the functions t in $S \rightarrow \overline{S}_\$$ where $S_\$$ is just S with a special “payoff” state $\$$ adjoined. The transitions in $\mathcal{R}.S$ give the probability of passage from initial s to final (proper) s' as $t.s.s'$; any deficit $1 - \sum_{s'} t.s.s'$ is interpreted as the probability of an immediate halt with payoff $t.s.\$ / (1 - \sum_{s'.S} t.s.s')$.

This formulation of the payoff has three desirable properties. The first is that the probabilistically *expected* halt-and-payoff is just $t.s.\$$. The second property is that we can consider the probabilities of outcomes from s to sum to 1 exactly (rather than no more than 1), since any deficit is “soaked up” in the probability of transit to payoff; that simplifies our operational interpretation.

The third property is that transitions preserve 1-boundedness of the expectations. Define the set of expectations $\mathcal{E}.S$ (over S) to be the set of 1-bounded functions $S \rightarrow [0, 1]$. If A in $\mathcal{E}.S$ gives a “post-expectation” $A.s'$ expected to be realised at state s' after transition t , then the “pre-expectation” at s before transition t is

$$t.s.\$ + \int_{t.s} A , \quad \text{where } t.s \text{ under } \int \text{ is restricted to states in } S \text{ proper.}^6$$

It is the expected value realised by making transition t from s to s' or possibly $\$$, and taking $A.s'$ in the former case. That this pre-expectation is also in $\mathcal{E}.S$ allows us to confine our work to the real interval $[0, 1]$ throughout.

Hence computation trees can be constructed by “pasting together” applications of transitions t_0, t_1, \dots drawn from $\mathcal{R}.S$, with branches to $\$$ being tips. The probabilities attached to the individual steps then generate a distribution over computational paths (which is defined by the sigma-algebra of extensions of finite sequences, a well-known construction [10]).

We use the relation \leq — “everywhere no more than” between expectations (thus replacing “implies”):

$$A \leq A' \quad \text{iff} \quad (\forall s: S \cdot A.s \leq A'.s) .$$

In our interpretations we will use *valuations* in the usual way. Given a formula ϕ , a valuation \mathcal{V} does four things: (i) it maps each A in ϕ to a fixed expectation in $\mathcal{E}.S$; (ii) it maps each occurrence of K to a fixed finite set of probabilistic

⁴ The restriction to the positive fragment is for the usual reason: that the interpretation of any expression $(\lambda X \cdot \phi)$, constructed according to the given rules, should yield a monotone function of X .

⁵ They correspond to the “game rounds” of Everett [8].

⁶ To avoid clutter we will assume this restriction where necessary in the sequel.

transitions in $\mathcal{R}.S$; (iii) it maps each occurrence of G to a predicate over S ; and (iv) it keeps track of the current instances of “unfoldings” of fixed points, by including mappings for bound variables X . (For notational economy, in (iv) we are allowing \mathcal{V} to take over the role usually given to a separate “environment” parameter.)

We make one simplification to our language, without compromising expressivity. The valuation \mathcal{V} assigns finite sets to all occurrences of K , hence we can replace each modality $\langle K \rangle \phi$ ($[K] \phi$) by an explicit maxjunct $\sqcup_{k:K} \{k\} \phi$ (minjunct $\sqcap_{k:K} \{k\} \phi$) of (symbols k denoting) transitions k in the set (denoted by) K . We do this because our interpretations do not distinguish between $\langle K \rangle$ or $[K]$ when K is a singleton set.

In the rest of this paper we shall therefore use the *reduced language* given by

$$\phi \hat{=} X \mid A \mid \{k\} \phi \mid \phi_1 \sqcap \phi_2 \mid \phi_1 \sqcup \phi_2 \mid \phi_1 \triangleleft G \triangleright \phi_2 \mid (\mu X \cdot \phi) \mid (\nu X \cdot \phi) .$$

We replace (ii) above in respect of \mathcal{V} by: (ii') it maps each occurrence of $\{k\}$ to a probabilistic transition in $\mathcal{R}.S$.

3 Logical interpretation: $qM\mu$ generalises Kozen's logic

In this section we recall how the quantitative logic for nondeterministic/probabilistic programs [14, 24] — from which we inherit the use of expectations, and the definition of $\|\{k\}\phi\|$ below — leads to a generalisation of Kozen's logical interpretation of μ -calculus for probabilistic transition systems.

Let ϕ be a formula and \mathcal{V} a valuation. We write $\|\phi\|_{\mathcal{V}}$ for its meaning, an expectation in $\mathcal{E}.S$ determined by the following rules:

1. $\|X\|_{\mathcal{V}} \hat{=} \mathcal{V}.X$.
2. $\|A\|_{\mathcal{V}} \hat{=} \mathcal{V}.A$.
3. $\|\{k\}\phi\|_{\mathcal{V}.s} \hat{=} \mathcal{V}.k.s.\$ + \int \|\phi\|_{\mathcal{V}}$.
4. $\|\phi' \sqcap \phi''\|_{\mathcal{V}.s} \hat{=} \|\phi'\|_{\mathcal{V}.s} \sqcap \|\phi''\|_{\mathcal{V}.s}$; and
 $\|\phi' \sqcup \phi''\|_{\mathcal{V}.s} \hat{=} \|\phi'\|_{\mathcal{V}.s} \sqcup \|\phi''\|_{\mathcal{V}.s}$.
5. $\|\phi' \triangleleft G \triangleright \phi''\|_{\mathcal{V}.s} \hat{=} \|\phi'\|_{\mathcal{V}.s}$ if ($\mathcal{V}.G.s$) else $\|\phi''\|_{\mathcal{V}.s}$.
6. $\|(\mu X \cdot \phi)\|_{\mathcal{V}} \hat{=} (\mu x \cdot \|\phi\|_{\mathcal{V}[x/X]})$ where, in the semantics on the *rhs*, by ($\mu x \cdot exp$) we mean the least fixed-point of the function $(\lambda x \cdot exp)$.
7. $\|(\nu X \cdot \phi)\|_{\mathcal{V}} \hat{=} (\nu x \cdot \|\phi\|_{\mathcal{V}[x/X]})$.

Note that in the valuation $\mathcal{V}[x/X]$, the variable X is mapped to the expectation x .

Lemma 1. The quantitative logic $qM\mu$ is well-defined — *For any ϕ in the language, and valuation \mathcal{V} , the interpretation $\|\phi\|_{\mathcal{V}}$ is a well-defined expectation in $\mathcal{E}.S$.*

Proof. Structural induction, arithmetic and that $(\mathcal{E}.S, \leq)$ is a complete partial order [19, 20].

4 Operational interpretation: $qM\mu$ generalises Stirling's game

In this section we give an alternative account of formulae ϕ (of the reduced language), in terms of a generalisation of Stirling's turn-based game [30]. It is played between two players, whom we refer to as respectively *Max* and *Min*. As in Sec. 3, we assume a probabilistic transition system $\mathcal{R}.S$ and a valuation \mathcal{V} . The game progresses through a sequence of *game positions*, each of which is either a pair (ϕ, s) where ϕ is a formula and s is a state in S , or a single (p) for some non-negative real number p representing a payoff. Following Stirling, we will use the idea of “colours” to represent place-holders for possible return to a fixed point. A sequence of game positions is called a *game path* and is of the form $(\phi_0, s_0), (\phi_1, s_1), \dots$ with (if finite) a payoff position (p_n) at the end. The initial formula ϕ_0 is the given ϕ , and s_0 is an *initial* state in S . A move from position (ϕ_i, s_i) to (ϕ_{i+1}, s_{i+1}) or to (p) is specified by the following rules.

1. If ϕ_i is $\phi \sqcap \phi'$ (resp. $\phi \sqcup \phi'$) then *Min* (*Max*) chooses one of the minjuncts (maxjuncts): the next game position is (ϕ'', s_i) , where ϕ'' is the chosen 'junct ϕ or ϕ' .
2. If ϕ_i is $\phi \triangleleft G \triangleright \phi'$, the next game position is (ϕ, s_i) if $\mathcal{V}.G.s_i$ holds, and otherwise it is (ϕ', s_i) .
3. If ϕ_i is $\{k\}\phi$ then the distribution $\mathcal{V}.k.s_i$ is used to choose either the next state s' in S or possibly the payoff state $\$$. If a state s' is chosen, then the next game position is (ϕ, s') ; if $\$$ is chosen, then the next position is (p) and is final, where p is the payoff $\mathcal{V}.k.s.\$ / (1 - \sum_{s' \in S} \mathcal{V}.k.s.s')$ ⁷ and the game terminates.
4. If ϕ_i is $\nu X \cdot \phi$ ($\mu X \cdot \phi$) then a fresh “colour” C is chosen from some infinite supply, and is bound to the formula $\phi[C/X]$ for later use; the next game position is (C, s_i) .⁸
5. If ϕ_i is C the next game position is (ϕ', s_i) , where ϕ' is the formula bound previously to C .
6. If ϕ_i is A then the game terminates in position (p) where $p = \mathcal{V}.A.s_i$.

A game path is said to be *valid* if it can occur as a sequence according to the above rules. Note that along any game path at most one colour can appear infinitely often:

Lemma 2. *All valid game paths are either finite, terminating at some payoff (p) , or infinite; if infinite, then exactly one colour appears infinitely often.*

Proof. *Stirling [30].*

⁷ Although p can exceed 1, if for example the denominator $1 - \sum_{s'} \mathcal{V}.k.s.s'$ is very small, the *expected* payoff for any game still lies between 0 and 1. If p is infinite, however, then the probability of that branch is 0 and it cannot be selected.

⁸ This use of *colours* is taken from Stirling [30]. The device allows easy determination, later on, of which recursion operator actually “caused” an infinite path, and is why μ and ν need not be distinguished at this point; they generate the same tree.

To complete the description of the game, one would normally give the winning/losing conditions. Here however we are operating over real- rather than Boolean values, and we speak of the “value” of the game. In the choices $\phi \sqcup \phi'$ ($\phi \sqcap \phi'$) player *Max*(*Min*) tries to maximise (minimise) a real-valued “payoff” associated with the game,⁹ defined as follows. There are a number of cases:

- The play is finite, terminating in a game state (p); in this case the payoff is p .
- The play is infinite and there is a colour C appearing infinitely often that was generated by a ν ; in this case the payoff is 1.
- The play is infinite and there is a colour C appearing infinitely often that was generated by a μ ; in this case the payoff is 0.

5 Example: strategic software development

Typical properties of probabilistic systems are usually cost-based. The following problem is an example — it also is an illustration of a property that lies strictly outside the scope of threshold-based probabilistic logics.

A Major company (M) has discovered a market opportunity for new software: and it could publish a low-quality version early, to capture that market and lock in its users; or it could publish later and at higher quality, but then running the risk that an Alternative supplier (A) will beat it to market. M 's marketing strategists have to determine the timing of their product launch.

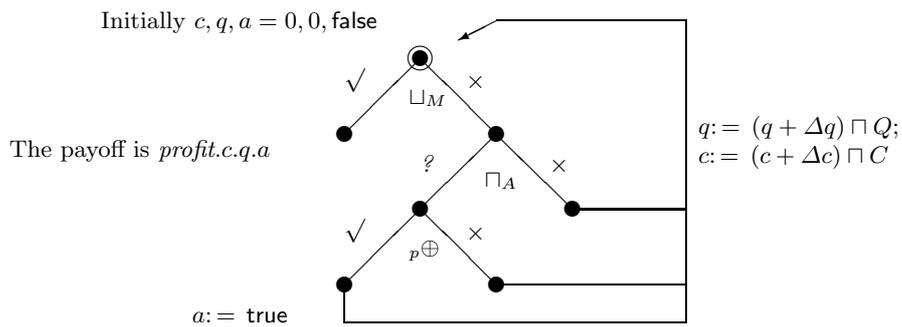
The marketeers' pressure for early publication comes from two sources — they know that the internal cost of the product (a bounded integer-valued variable c) will rise the longer it's left with the developers, and there is the increasing risk that A will publish first. They estimate that A will publish (recorded by Boolean variable a) with probability no more than some constant p per unit time.

The complementary argument, for delaying publication, is that the more time spent on development the better the quality (variable q) of the final product will be. That will save on maintenance costs later and will enhance the company's reputation.

Finally, because M is a much bigger company than A , it can be sure that if M publishes first then A won't publish at all. When M does publish, M 's profit is given by a function *profit.c.q.a* into $[0, 1]$ of the cost (c), the quality (q) and whether A has published already (a).

The situation is summed up by the transition system set out in Fig. 1: at each time step M can choose either to publish immediately (by selecting the left-hand branch at \sqcup_M , and terminating), or to postpone publication and continue development for another time step (by selecting the right-hand branch at \sqcup_M).

⁹ In fact attributing the wins/losses to the two players makes it into a *zero-sum* game.



If repeated forever, then the payoff is zero.

- \sqcup_M — Angelic choice (\exists modality) of whether M publishes (\checkmark) or not (\times). (Chosen by Max in the game.)
- \sqcap_A — Demonic choice (\forall modality) of whether A even considers ($?$) publishing. (Chosen by Min in the game.)
- $p \oplus$ — Probabilistic choice of whether A publishes. (“Chosen” by *chance*.)

Function $\text{profit}.c.q.a$ determines the payoff realised by M from publication of a product with quality q and development costs c , and depends on whether A has published or not (Boolean a).

In general, given a definition of $\text{profit}.c.q.a$, company M 's optimal payoff can be computed to reveal the best timing for the launch of their product.

Fig. 1. Strategic software development

If M chooses the latter option then it risks A 's publishing first: in the worst case (from M 's point of view, the left-hand branch at \sqcap_A) A does so with probability p (the left-hand probabilistic branch at ${}_p\oplus$). These steps are repeated until M publishes or — the non-terminating case — forever, if M never publishes at all. In that latter case, the payoff is 0.

The utility of our game interpretation in Sec. 4 is that we can easily use the intuition it provides to write a formula describing the above system over the state space (c, q, a) : the expected payoff is $(\mu X \cdot \text{profit} \sqcup [\mathbf{K}]X)$, where \mathbf{K} denotes the set $\{k_0, k_1\}$ of transitions

$$\begin{aligned} k_0 &\hat{=} q := (q + \Delta q) \sqcap Q; & c &:= (c + \Delta c) \sqcap C \\ k_1 &\hat{=} (a := \text{true} \; {}_p\oplus \; a := \text{false}); & & k_0, \end{aligned}$$

and where “;” is sequential composition in the small programming language we are using to describe the transitions. (We assume a valuation mapping *profit* to *profit* etc.)

The choice implicit in $[\mathbf{K}]$ is the way we express probability ranges, if the problem demands it: here, it is that the probability of A 's publishing is “no more than p ” is coded up as a choice between exactly p (if k_1 is chosen at every play) and 0 (if k_0 is always chosen). Since the choice \sqcap_A can be resolved by player *Min* to any probability in $[0, 1]$, the overall probability range for A 's publishing is $[0, p]$.

Note that our transitions k_0, k_1 take an initial state to a full- (rather than sub-) distribution over final states, and that in this example there is no “immediate payoff” component (i.e. it is 0). This is of course a special case of the transitions we allow in $\mathcal{R.S}$: in the notation of the introduction we just have $t.s.\$ = 0$ and $\sum_{s':S} t.s.s' = 1$. The role of (and need for) the extra generality is discussed in the conclusion.

We have used the least (rather than greatest) fixed point because “never publishing” pays 0.

In the *reduced language* (at end of Sec. 2) we would write our formula as

$$\text{Game} \hat{=} (\mu X \cdot \text{profit} \sqcup (\{k_0\}X \sqcap \{k_1\}X)) ,$$

and then the Kozen-style interpretation $\| \text{Game} \|_{\mathcal{V}.s_0}$, given appropriate values *profit*, k_0, k_1 supplied by \mathcal{V} , is a well-defined expectation as set out in Sec. 3. For example when p is $1/3$, and *profit.c.q.a* is defined to be¹⁰ (0 if a else $q - c$) — a simple definition assuming all payoff is lost if A reaches the market first — then a short calculation gives the value $(8/9)(\Delta q - \Delta c)$, assuming $\Delta q \geq \Delta c$, initial state $q, c, a = 0, 0, \text{false}$, and that the bounds Q, C are not too low. As we argue for the general case in Sec. 6, this turns out to be the same as M 's optimal expected payoff in the game, whatever A 's strategy might be.

Alternatively, in the Stirling-style interpretation we generate a game-tree by “unfolding” the transition system in Fig. 1. At each unfolding, M and A need

¹⁰ The profit function should be 1-bounded, but to avoid clutter we have not scaled it down here (e.g. by dividing by $Q \sqcup C$).

to select a branch; and their selections could be different each time they revisit their respective decision points. Let σ_M and σ_A be sequences (possibly infinite) of choices to be made by M and A . When they follow those sequences, the resulting game-tree generates a well-defined probability distribution over valid game paths [10]. Anticipating the next section, let $\llbracket \phi \rrbracket_{\mathcal{V}}^{\sigma_M, \sigma_A}$ denote that path distribution: we can now describe M 's actual payoff — a function of those strategies — as an expected value

$$P.\sigma_M.\sigma_A \hat{=} \int_{\llbracket \phi \rrbracket_{\mathcal{V}}^{\sigma_M, \sigma_A}} \text{“profit applied to the final state”} ,$$

with the understanding that the random variable over paths, in the integral's body, yields 0 if in fact there is no final state (infinite path).

As usual in game theory, when the actual strategies of the two players are unknown (as in this case), we define the *value* of the game to be the the minimax over all strategy sequences of the expected payoff — but the minimax is well-defined only when it is the same as the maximin, that is only if

$$\sqcup_{\sigma_M} \sqcap_{\sigma_A} P.\sigma_M.\sigma_A = \sqcap_{\sigma_A} \sqcup_{\sigma_M} P.\sigma_M.\sigma_A .$$

In some cases, the value of a game can be realised by *memoriless strategies* — roughly speaking, a memoriless strategy is independent of the number of unfoldings of the game-tree. Memoriless strategies are particularly important for the efficient computation of expected payoffs [9], and in Sec. 6 we show they suffice for $qM\mu$ when the state space is finite.

To summarise, in the next section we show that the techniques used in this example are valid in general — that is, that the values of the games set out in Sec. 4 are all well-defined, that they can be realised by memoriless strategies if the state space is finite, and that the value corresponds exactly to the denotational interpretation of Sec. 3.

For the current example, those results justify our using the Kozen-style interpretation to calculate M 's optimal profit in the game, which in this simple case led to a direct calculation. For more complex formulae, the optimal payoff is determined by constructing explicit automata based on the corresponding games, and then applying model-checking methods [2].

6 Proof of duality

In this section we give our main result, the equivalence of the two interpretations of a $qM\mu$ formula: the operational, “Stirling-game” interpretation, and the denotational “Kozen-logic” interpretation. In both cases we must address explicitly the question of *strategies*, and whether they can or cannot have “memory” of where the game or transition system has gone so far.

6.1 Duality over given strategies

We begin with the Stirling interpretation, and our first step will be to explain how the games can be formalised provided the players' strategies are decided beforehand.

The current position of a game — as we saw in Sec. 4 — is a formula/state pair. We introduce two *strategy functions* called $\underline{\sigma}$ and $\bar{\sigma}$, to formalise the players' decisions as they go along: the functions are of type “finite-game-path to Boolean”, and the player *Min* (resp. *Max*), instead of deciding “on the fly” how to interpret a decision point \sqcap (resp. \sqcup), takes the strategy function $\underline{\sigma}$ (resp. $\bar{\sigma}$) and applies that to the sequence of game positions traversed so far with, say, result “true” meaning “take the left subformula”.

These strategies model full memory, because each is given as an argument the complete history of the game up to its point of use. (Note that the history includes the current state s .)

The formalisation of the Stirling game is then in two stages. In the first stage we construct a (possibly infinite) probabilistically-branching game-tree $\llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}}.s$, using the given formula ϕ , the initial state s and the pre-packaged strategy functions $\underline{\sigma}, \bar{\sigma}$.

For the second stage we use a function *Val*, from valid game paths to the non-negative reals, which is defined to give exactly the “payoff” described at the end of Sec. 4. Then we have

Definition 1. Value of fixed-strategy Stirling game — *The value of a game played from formula ϕ and initial state s , with strategies $\underline{\sigma}, \bar{\sigma}$, is given by the expected value*

$$\int_{\llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}}.s} \text{Val}$$

of Val over the game-tree $\llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}}.s$ generated by the strategies and the initial state.

The argument that this is well defined is the usual one, based on showing that Val is a measurable function over the sigma-algebra defined by the tree.

Our second step is to show that the above game corresponds to a Kozen-style interpretation over the same data: that is, we augment the semantics of Sec. 3 with the same strategy functions. For clarity we use slightly different brackets $\llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}}$ for the extended semantics.

The necessary alterations to the rules in Sec. 3 are straightforward, the principal one being that in Case 4, instead of taking a minimum or maximum, we use the argument $\underline{\sigma}$ or $\bar{\sigma}$ as appropriate to determine whether to carry on with ϕ' or with ϕ'' . (A technical complication is then that all the definitions have to be changed so that the “game sequence so far” is available to $\underline{\sigma}$ and $\bar{\sigma}$ when required. That can be arranged for example by introducing an extra “path so far” argument and passing it, suitably extended, on every right-hand side.)

We then have our first equivalence:

Lemma 3. Equivalence of given-strategy games and logic — For all $qM\mu$ formulae ϕ , valuations \mathcal{V} , states s and strategies $\underline{\sigma}, \bar{\sigma}$, we have

$$\int_{\llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}}.s} Val = \llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}}.s .$$

Proof. (sketch) The proof is by structural induction over ϕ , straightforward except when least- or greatest fixed-points generate infinite trees. In those cases we consider longer and longer finite sub-trees of the infinite tree: as well as the valid paths already described, they may contain extra finite paths, ending in μ - or ν -generated colours that indicate unfoldings yet to come.

Extend the path-valuation function Val so that it assigns 0 (resp. 1) to a finite path ending in a μ - (resp. ν -) colour. It can be shown that the function $\int Val$ is continuous over a limit of a subtree-ordered sequence of partial trees, ^(.) provided there is a single colour such that every extra finite path in any tree in the sequence terminates in that colour. In that case the game's overall value is the limit of the non-decreasing (resp. non-increasing) sequence of values assigned by the extended \mathcal{V} to the finite trees.

Since, in a single μ - (resp. ν -) structural induction step, the infinite paths assigned 0 (resp. 1) by Val are exactly those containing infinitely many occurrences of the associated fixed colour, a limit of trees as above can be constructed. And each of any μ - (resp. ν -) generated infinite path's prefixes is assigned 0 (resp. 1) appropriately by the above extension of Val to finite colour-terminated paths.

For a full proof see [26].

Lem. 3 is the key to completing the argument that the value of the Stirling game is the minimax over all strategies of the expected payoff: recalling the issues raised at the end of Sec. 5, we must do that to show it to be well defined. That is, in the notation of this section we must establish

$$\sqcap_{\underline{\sigma}} \sqcup_{\bar{\sigma}} \int_{\llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}}.s} Val = \sqcup_{\bar{\sigma}} \sqcap_{\underline{\sigma}} \int_{\llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}}.s} Val . \quad (1)$$

The utility of Lem. 3 is that it allows us to carry out that argument in a denotational rather than operational context — we can avoid the integrals, and simply show $\sqcap_{\underline{\sigma}} \sqcup_{\bar{\sigma}} \llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}} = \sqcup_{\bar{\sigma}} \sqcap_{\underline{\sigma}} \llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}}$ instead.

In fact we show (1) to be even simpler — it is given by the original Kozen interpretation $\llbracket \phi \rrbracket_{\mathcal{V}}$ with its \sqcap and \sqcup operators still in place, which therefore is the value of the Stirling game.

6.2 Full duality, and memoriless strategies

Let formula $\phi_{\underline{G}}$ be derived from ϕ by replacing each operator \sqcap in ϕ by a specific state predicate drawn from a tuple \underline{G} of our choice, possibly a different predicate for each syntactic occurrence of \sqcap . Similarly we write $\phi_{\bar{G}}$ for the derived formula

in which all instances of \sqcup are replaced one-by-one by successive state predicates in a tuple $\overline{\mathbf{G}}$. With those conventions, we will appeal to Lem. 10 that for all $qM\mu$ formulae ϕ over a finite¹¹ state space S , and valuations \mathcal{V} , there exist predicate (tuples) $\underline{\mathbf{G}}$ and $\overline{\mathbf{G}}$ as above such that $\|\phi_{\underline{\mathbf{G}}}\|_{\mathcal{V}} = \|\phi\|_{\mathcal{V}} = \|\phi_{\overline{\mathbf{G}}}\|_{\mathcal{V}}$.

For example, if the formula ϕ is $(\mu X \cdot A_1 \sqcup (\nu Y \cdot A_2 \sqcap \{k\} (A_3 \sqcup (X \triangleleft G \triangleright Y))))$, then we are saying we can find tuples $\underline{\mathbf{G}} \hat{=} (\underline{\mathbf{G}}_1)$ and $\overline{\mathbf{G}} \hat{=} (\overline{\mathbf{G}}_1, \overline{\mathbf{G}}_2)$ so that

$$\begin{aligned} \phi_{\underline{\mathbf{G}}} &\hat{=} (\mu X \cdot A_1 \sqcup (\nu Y \cdot A_2 \triangleleft \underline{\mathbf{G}}_1 \triangleright \{k\} (A_3 \sqcup (X \triangleleft G \triangleright Y)))) && \text{and} \\ \phi_{\overline{\mathbf{G}}} &\hat{=} (\mu X \cdot A_1 \triangleleft \overline{\mathbf{G}}_1 \triangleright (\nu Y \cdot A_2 \sqcap \{k\} (A_3 \triangleleft \overline{\mathbf{G}}_2 \triangleright (X \triangleleft G \triangleright Y)))) \end{aligned}$$

are both equivalent to ϕ under $\|\cdot\|_{\mathcal{V}}$.¹²

The proof of Lem. 10 is by induction, intricate only in one case, where we rely on techniques due to Everett [8].¹³ That part of the proof, together with several preliminary lemmas, is given in the appendix; the full proof is given elsewhere [26].

Lemma 4. Minimax well-defined for Kozen interpretation — *For all $qM\mu$ formulae ϕ , valuations \mathcal{V} and strategies $\underline{\sigma}, \overline{\sigma}$, we have*

$$\sqcap_{\underline{\sigma}} \sqcup_{\overline{\sigma}} \|\phi\|_{\mathcal{V}}^{\underline{\sigma}, \overline{\sigma}} = \sqcup_{\overline{\sigma}} \sqcap_{\underline{\sigma}} \|\phi\|_{\mathcal{V}}^{\underline{\sigma}, \overline{\sigma}}. \quad (2)$$

Proof. From monotonicity, we need only prove $\text{lhs} \leq \text{rhs}$. Note that from Lem. 10 we have predicates $\overline{\mathbf{G}}$ and $\underline{\mathbf{G}}$ satisfying

$$\|\phi_{\underline{\mathbf{G}}}\|_{\mathcal{V}} = \|\phi\|_{\mathcal{V}} = \|\phi_{\overline{\mathbf{G}}}\|_{\mathcal{V}}, \quad (3)$$

a fact which we use further below.

To begin with, using the predicates $\underline{\mathbf{G}}$ from (3), we start from the lhs of (2) and observe that

$$\sqcap_{\underline{\sigma}} \sqcup_{\overline{\sigma}} \|\phi\|_{\mathcal{V}}^{\underline{\sigma}, \overline{\sigma}} \leq \sqcup_{\overline{\sigma}} \|\phi_{\underline{\mathbf{G}}}\|_{\mathcal{V}}^{\overline{\sigma}}, \quad (4)$$

(in which on the right we omit the now-ignored $\underline{\sigma}$ argument), because the $\sqcap_{\underline{\sigma}}$ can select exactly those predicates $\underline{\mathbf{G}}$ by an appropriate choice of $\underline{\sigma}$. We then eliminate the explicit strategies altogether by observing that

$$\sqcup_{\overline{\sigma}} \|\phi_{\underline{\mathbf{G}}}\|_{\mathcal{V}}^{\overline{\sigma}} \leq \|\phi_{\underline{\mathbf{G}}}\|_{\mathcal{V}}, \quad (5)$$

because the simpler $\|\cdot\|$ -style semantics on the right interprets \sqcup as maximum, which cannot be less than the result of appealing to some strategy function $\overline{\sigma}$.

We can now continue on our way towards the rhs of (2) as follows:

$$\begin{aligned} & \|\phi_{\underline{\mathbf{G}}}\|_{\mathcal{V}} && \text{carrying on from (5)} \\ = & \|\phi\|_{\mathcal{V}} && \text{first equality at (3)} \end{aligned}$$

¹¹ Finiteness is needed in Case $\overline{\mathbf{G}}$ of the lemma's proof.

¹² In fact it is easy to show that all three formulae are then equivalent to $\phi_{\underline{\mathbf{G}}, \overline{\sigma}}$, but we do not need that.

¹³ Unfortunately Everett's work as it stands is less than we need, so although we borrow his techniques we cannot simply appeal to his result as a whole.

$$\begin{aligned}
&= && \|\phi_{\overline{\mathbb{G}}}\|_{\mathcal{V}} && \text{second equality at (3)} \\
&\leq && \sqcap_{\underline{\sigma}} \|\phi_{\overline{\mathbb{G}}}\|_{\overline{\mathcal{V}}^{\underline{\sigma}}} && \text{as for (5) above, backwards} \\
&\leq && \sqcup_{\overline{\sigma}} \sqcap_{\underline{\sigma}} \|\phi\|_{\overline{\mathcal{V}}^{\underline{\sigma}, \overline{\sigma}}} && \text{as for (4) above, backwards}
\end{aligned}$$

and we are done.

The proof above establishes the duality we seek between the two interpretations.

Theorem 1. Value of Stirling game — *The value of a Stirling game is well-defined, and equals $\|\phi\|_{\mathcal{V}}$.*

Proof. Lem. 3 and Lem. 4 establish the equality (1), for well-definedness; the stated equality with $\|\phi\|_{\mathcal{V}}$ occurs during the proof of the latter.

Finally, we have an even tighter result about the players' strategies:

Lemma 5. Memoriless strategies — *There exists a memoriless strategy $\overline{\mathbb{G}}$ which, if followed by player Max, achieves the value of the Stirling game against all strategies of player Min. (A similar result holds for player Min.)*

Proof. Directly from Lem. 10 and Thm. 1.

7 Conclusion

Our main contribution has been to introduce a novel game-based interpretation for the quantitative μ -calculus $qM\mu$ over probabilistic/angelic/demonic transition systems, probabilistically generalising Stirling's game interpretation of the standard μ -calculus, and to show it equivalent to the existing Kozen-style interpretation of $qM\mu$. The interpretations are general enough to specify cost-based properties of probabilistic systems — and many such properties lie outside standard temporal logic. The Stirling-style interpretation is close to automata-based approaches, whilst the Kozen-style logic (studied more extensively elsewhere [20]) provides an attractive proof system.

Part of our generalisation has been to introduce the Everett-style “payoff states” $\$$ into Stirling's generalised games. Although many presentations of probabilistic transitions (including our earlier work) do not include the extra state, giving instead simply functions from S to \overline{S} which in effect take the primitive elements of formulae to be probabilistic programs, here our primitive elements are small probabilistic *games* [8]. The probabilistic programs are just the simpler special case of payoff 0. The full proof [26] of Lem. 10 makes that necessary, since we treat the $\underline{\mathbb{G}}/\nu$ case via a duality, appealing to the $\overline{\mathbb{G}}/\mu$ case. But it is a duality under which probabilistic programs are not closed, whereas the slightly more general probabilistic games are closed. Thus we have had to prove a slightly more general result.

8 Related work

Probabilistic temporal logics, interpreted over nondeterministic/probabilistic transition systems, have been studied extensively, most notably by de Alfaro [5], Jonsson [11], Segala [29] and Vardi [32]. Condon [3] considered the complexity of underlying transition systems like ours, including probabilistic (but $_{1/2}\oplus$ only), demonic and angelic choice, but without our more general expectations and payoffs. Monniaux [18] uses Kozen’s deterministic formulation together with demonic program inputs to analyse systems via abstract interpretation [4].

The quantitative μ -calculus has not received as much attention. Huth and Kwiatkowska [12] use real-valued expressions based on expectations, and they have investigated model-checking approaches to evaluating them; but they do not provide an operational interpretation of the logic, nor have they exploited its algebraic properties [20]. Narasimha et. al. [25] use probability thresholds, and restrict to the alternation-free fragment of the μ -calculus; however for that fragment they do provide an operational interpretation which selects the proportion of paths that satisfy the given formula. Their transition systems are deterministic.

De Alfaro and Majumdar [7] use $qM\mu$ to address an issue similar to, but not the same as ours: in the more general context of concurrent games, they show that for every LTL formula Ψ one can construct a $qM\mu$ formula ϕ such that $\|\phi\|_{\mathcal{V}}$ is the greatest assured probability that Player 1 can force the game path to satisfy Ψ .

The difference between de Alfaro’s approach and ours can be seen by considering the formula $\Psi \triangleq (\mu X \cdot \{k\}\text{atB} \sqcup \{k\}X)$ over the transition system

$$\mathcal{V}.k \triangleq (s := A \text{ }_{1/2}\oplus s := B) \text{ if } (s = A) \text{ else } (s := A)$$

operating on state space $\{A, B\}$. (In fact formula Ψ expresses the notorious $AFAX\text{atB}$ [31] in the temporal subset [20] of $qM\mu$, where $\mathcal{V}.\text{atB}.s \triangleq 1$ if $(s = B)$ else 0 . Player 1 can force satisfaction of Ψ with probability 1 in this game, since the only path for which it fails (all A ’s) occurs with probability 0; so de Alfaro’s construction yields a *different* formula ϕ such that $\|\phi\|_{\mathcal{V}} = 1$.

Yet $\|\Psi\|_{\mathcal{V}}$ for the original formula is only $1/2$, which is the value of the *Stirling* game played in this system. It is “at each step, seek to maximise (\sqcup) the payoff, depending on whether after the following step ($\{k\}$) you will accept atB and terminate, or go around again (X).” Note that the decision “whether to repeat after the next step” is made *before* that step is taken. (Deciding after the step would be described by the formula $(\mu X \cdot \{k\}(\text{atB} \sqcup X))$.) The optimal strategy for *Max* is of course given by $\Psi_{\text{atA}} \triangleq (\mu X \cdot \{k\}\text{atB} \text{ if } \text{atA} \text{ else } \{k\}X)$.

Finally, our result Lem. 5 for memoriless strategies holds for all $qM\mu$ formulae, whereas (we believe) de Alfaro et. al. treat only a subset, those formulae encoding the automata used in their construction.

References

1. M. Ben-Ari, A. Pnueli, and Z. Manna. The temporal logic of branching time. *Acta Informatica*, 20:207–226, 1983.
2. Andrea Bianco and Luca de Alfaro. Model checking of probabilistic and nondeterministic systems. In *Foundations of Software Technology and Theoretical Computer Science*, volume 1026 of *LNCS*, pages 499–512, December 1995.
3. A. Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203–224, 1992.
4. P. Cousot and R. Cousot. Abstract interpretation frameworks. *Journal of Logic and Computation*, 2(2):511–547, 1992.
5. Luca de Alfaro. Temporal logics for the specification of performance and reliability. In *STACS '97*, volume 1200 of *LNCS*, 1997.
6. Luca de Alfaro. Computing minimum and maximum reachability times in probabilistic systems. In *Proceedings of CONCUR '99*, LNCS. Springer Verlag, 1999.
7. Luca de Alfaro and Rupak Majumdar. Quantitative solution of omega-regular games. In *Proc. STOC '01*, 2001.
8. H. Everett. Recursive games. In *Contributions to the Theory of Games III*, volume 39 of *Ann. Math. Stud.*, pages 47–78. Princeton University Press, 1957.
9. J. Filar and O.J. Vrieze. *Competitive Markov Decision Processes — Theory, Algorithms, and Applications*. Springer Verlag, 1996.
10. G. Grimmett and D. Welsh. *Probability: an Introduction*. Oxford Science Publications, 1986.
11. Hans Hansson and Bengt Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6:512–535, 1994.
12. Michael Huth and Marta Kwiatkowska. Quantitative analysis and model checking. In *Proceedings of 12th annual IEEE Symposium on Logic in Computer Science*, 1997.
13. D. Kozen. Semantics of probabilistic programs. *Journal of Computer and System Sciences*, 22:328–350, 1981.
14. D. Kozen. A probabilistic PDL. In *Proceedings of the 15th ACM Symposium on Theory of Computing*, New York, 1983. ACM.
15. D. Kozen. Results on the propositional μ -calculus. *Theoretical Computer Science*, 27:333–354, 1983.
16. A.K. McIver and C. Morgan. Demonic, angelic and unbounded probabilistic choices in sequential programs. *Acta Informatica*, 37:329–354, 2001.
17. A.K. McIver, C.C. Morgan, and J.W. Sanders. Probably Hoare? Hoare probably! In A.W. Roscoe, editor, *A Classical Mind: Essays in Honour of CAR Hoare*. Prentice-Hall, 1999.
18. David Monniaux. Abstract interpretation of probabilistic semantics. In *International Static Analysis Symposium (SAS '00)*, volume 1824 of *LNCS*. Springer Verlag, 2000.
19. Carroll Morgan and Annabelle McIver. A probabilistic temporal calculus based on expectations. In Lindsay Groves and Steve Reeves, editors, *Proc. Formal Methods Pacific '97*. Springer Verlag Singapore, July 1997. Available at [27].
20. Carroll Morgan and Annabelle McIver. An expectation-based model for probabilistic temporal logic. *Logic Journal of the IGPL*, 7(6):779–804, 1999. Also available via [27].
21. Carroll Morgan and Annabelle McIver. *pGCL*: Formal reasoning for random algorithms. *South African Computer Journal*, 22, March 1999. Also available at [27].

22. Carroll Morgan and Annabelle McIver. Almost-certain eventualities and abstract probabilities in the quantitative temporal logic qTL . In *Proceedings CATS '01*. Elsevier, 2000. Also available at [27]; to appear in *Theoretical Computer Science*.
23. C.C. Morgan and A.K. McIver. Cost analysis of games using program logic. In *Proc. of the 8th Asia-Pacific Software Engineering Conference (APSEC 2001)*, December 2001. Abstract only: full text available at [27].
24. C.C. Morgan, A.K. McIver, and K. Seidel. Probabilistic predicate transformers. *ACM Transactions on Programming Languages and Systems*, 18(3):325–353, May 1996.
25. N. Narasimha, R. Cleaveland, and P. Iyer. Probabilistic temporal logics via the modal mu-calculus. In *Proceedings of the Foundation of Software Sciences and Computation Structures, Amsterdam*, number 1578 in LNCS, pages 288–305, 1999.
26. Draft presentations of the full proofs of these results can be found via entry *Games02* at the web site [27].
27. PSG. Probabilistic Systems Group: Collected reports. <http://web.comlab.ox.ac.uk/oucl/research/areas/probs/bibliography.html>.
28. M.O. Rabin. The choice-coordination problem. *Acta Informatica*, 17(2):121–134, June 1982.
29. Roberto Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, MIT, 1995.
30. Colin Stirling. Local model checking games. In *CONCUR '95*, volume 962 of LNCS, pages 1–11. Springer Verlag, 1995. Extended abstract.
31. M. Y. Vardi. Branching vs. linear time: Final showdown. In *Seventh International Conference on Tools and Analysis of Systems, Genova*, number 2031 in LNCS, April 2001.
32. Moshe Y. Vardi. A temporal fixpoint calculus. In *Proc. 15th Ann. ACM Symp. on Principles of Programming Languages*. ACM, January 1988. Extended abstract.
33. I. Walukiewicz. Notes on the propositional mu-calculus: Completeness and related results. Technical Report BRICS NS-95-1, BRICS, Dept. Comp. Sci., University Aarhus, 1995. Available at <http://www.brics.aau.dk/BRICS/>.

A Memoriless strategies suffice

We show that for any formula ϕ , possibly including \Box and \sqcup strategy operators, there are specific state predicates (collected into tuples $\underline{\mathbb{G}}$ and $\overline{\mathbb{G}}$) that can replace the strategy operators without affecting the value of the formula. The inductive proof is straightforward except for replacement of \sqcup within μ (and, dually, replacement of \Box within ν). For this $\overline{\mathbb{G}}/\mu$ case we need several technical lemmas and definitions; the other cases are set out at [26].

Because the argument in this section is mainly over properties of real-valued functions, we shift to a more mathematical style of presentation. Variables f, g, \dots denote (Curried) functions of type expectation(s) to expectation, and w, x, \dots are expectations in $\mathcal{E}.S$. For function f of one argument we write $\mu.f$ for its least fixed-point.

Definition 2. Almost-linear — *Say that an expectation-valued function f of possibly several expectation arguments x, y, \dots, z is almost-linear if it can be written in the form*

$$f.x.y \cdots z \hat{=} w + g.x + h.y + \cdots + i.z, \quad (6)$$

where w is an expectation and g, h, \dots, i are linear expectation-valued functions of their single arguments.

Lemma 6. Every \sqcap/\sqcup -free formula ϕ , possibly containing free expectation variables X, Y, \dots, Z , denotes an almost-linear function of the values assigned to those arguments.

Proof. (sketch) What we are claiming is that the function

$$f.x.y \cdots .z \quad \hat{=} \quad \|\phi\|_{\mathcal{V}[x,y \cdots z/X,Y \cdots Z]}$$

can be written in the form given on the right at (6), provided ϕ contains no \sqcap or \sqcup . This is a straightforward structural induction over ϕ , given in full at [26].

Definition 3. Almost less-than — For non-negative reals a, b , write $a \ll b$ for $a > 0 \Rightarrow a < b$; write the same for the pointwise-extended relation over expectations. Note that $a \ll b$ implies $a \leq b$ on this domain.

Definition 4. ok functions — Say that an expectation-to-expectation function f of one argument is ok if for all expectations x with $x \ll f.x$ we have that $(\sqcup n \cdot f^n.x) = \mu.f$.

(Note that $(\sqcup n \cdot f^n.x) \geq (\sqcup n \cdot f^n.0) = \mu.f$ trivially if f is monotonic.)

Lemma 7. If f is almost-linear then f is ok in each argument separately.

Proof. (sketch) This technical result says in effect that if expectation x is increased by applying an almost-linear f of one argument, then the states s at which $x.s$ is non-zero eventually do not affect the terms $f^n.x$ in the supremum — whence we might as well take $f^n.0$, which limit gives the right-hand side.

A full proof is given at [26].

Lemma 8. All \sqcap/\sqcup -free formulae ϕ denote ok functions of their free expectations X, Y, \dots, Z taken separately.

Proof. Lemmas 6 and 7.

Lemma 9. For any monotonic and continuous¹⁴ function f over expectations, and any $\varepsilon > 0$, there is an expectation x such that $\mu.f - \varepsilon \leq x$ and $x \ll f.x$, where we write “ $-\varepsilon$ ” to subtract the constant function of that value.

This technical result forms the core of Everett’s argument [8]; note it does not depend on f ’s being ok. We give a full proof of it at [26].

We can now sketch the proof of the main result of this section.

Lemma 10. Fixed strategies suffice — For any formula ϕ , possibly containing strategy operators \sqcap/\sqcup , and valuation \mathcal{V} , there are state-predicate tuples $\underline{G}/\overline{G}$ — possibly depending on \mathcal{V} — such that

$$\|\phi_{\underline{G}}\|_{\mathcal{V}} = \|\phi\|_{\mathcal{V}} = \|\phi_{\overline{G}}\|_{\mathcal{V}} .$$

¹⁴ This is continuity in the usual sense in analysis.

Proof. (sketch) We give only the μ -case of an otherwise straightforward induction over the size of ϕ ; a full proof may be found at [26].

Suppose we are considering $(\mu X \cdot \phi)$. Let f be the function denoted by ϕ with respect to a single expectation-valued argument x supplied for the variable X , with the values of any other free variables in ϕ fixed by the environment \mathcal{V} ; for any $\underline{G}, \overline{G}$ let functions $f_{\underline{G}}$ and $f_{\overline{G}}$ s be derived similarly from $\phi_{\underline{G}}$ and $\phi_{\overline{G}}$.

Case \underline{G} : We must show $\mu.f = \mu.f_{\underline{G}}$ for some \underline{G} .¹⁵ note that $\mu.f \leq \mu.f_{\underline{G}}$ trivially, since $f \leq f_{\underline{G}}$. Since ϕ is smaller in size than $(\mu X \cdot \phi)$, our inductive hypothesis provides for any x a \underline{G}_x so that $f_{\underline{G}_x}.x = f.x$; take $x = \mu.f$ and therefore choose \underline{G} so that $f_{\underline{G}}.(\mu.f) = f.(\mu.f) = \mu.f$. Thus $\mu.f$ is a fixed-point of $f_{\underline{G}}$, whence immediately $\mu.f_{\underline{G}} \leq \mu.f$.

Case \overline{G} : In this case must show $\mu.f = \mu.f_{\overline{G}}$ for some \overline{G} ; again it is trivial that $\mu.f_{\overline{G}} \leq \mu.f$ for any \overline{G} .

For the other direction, in fact we show that for any $\varepsilon > 0$ there is a \overline{G}_ε such that $\mu.f_{\overline{G}_\varepsilon} \geq \mu.f - \varepsilon$ — whence the existence of a single \overline{G} satisfying $\mu.f_{\overline{G}} \geq \mu.f$ follows from the finiteness of the state space (since the set of possible strategy tuples for this f is therefore finite as well, and so there must be one that works for all ε).

Because we know inductively that f is a minimax¹⁶ over strategy tuples $\underline{G}', \overline{G}'$ of almost-linear functions $f_{\underline{G}', \overline{G}'}$, that those functions are continuous by construction, and that the minimax is finite because there are only finitely many strategy tuples $\underline{G}', \overline{G}'$ for this f , we know that f is continuous itself, and by Lem. 9 we therefore have an expectation x_ε with

$$\mu.f - \varepsilon \leq x_\varepsilon \quad \text{and} \quad x_\varepsilon \ll f.x_\varepsilon. \quad (7)$$

To get our result we need only show in addition that $x_\varepsilon \leq \mu.f_{\overline{G}_\varepsilon}$ for some \overline{G}_ε .

From our inductive hypothesis we can choose \overline{G}_ε so that $f.x_\varepsilon = f_{\overline{G}_\varepsilon}.x_\varepsilon$, whence from (7) we have $x_\varepsilon \ll f_{\overline{G}_\varepsilon}.x_\varepsilon$.

Now in fact $f_{\overline{G}_\varepsilon}$ is ok (see below), so that from Def. 4 we have $(\sqcup n \cdot f_{\overline{G}_\varepsilon}^n.x_\varepsilon) = \mu.f_{\overline{G}_\varepsilon}$; but then $x_\varepsilon \ll f_{\overline{G}_\varepsilon}.x_\varepsilon \leq (\sqcup n \cdot f_{\overline{G}_\varepsilon}^n.x_\varepsilon) = \mu.f_{\overline{G}_\varepsilon}$. Thus we have $x_\varepsilon \leq \mu.f_{\overline{G}_\varepsilon}$, and we are done.

To see that $f_{\overline{G}_\varepsilon}$ is ok, consider any x such that $x \ll f_{\overline{G}_\varepsilon}.x$. We note that the formula $\phi_{\overline{G}_\varepsilon}$ — from which $f_{\overline{G}_\varepsilon}$ is derived — is the same size as ϕ , and therefore smaller than $(\mu X \cdot \phi)$,¹⁷ so that it satisfies the inductive hypothesis and we can apply the argument of Case \underline{G} to it: thus we have a \underline{G}' with $\mu.f_{\overline{G}_\varepsilon} = \mu.f_{\underline{G}', \overline{G}_\varepsilon}$.

¹⁵ Note that $(\mu X \cdot \phi)_{\underline{G}}$ is the same as $(\mu X \cdot \phi_{\underline{G}})$ — it is syntactic substitution — so that $\mu.(f_{\underline{G}})$ is indeed the correct denotation.

¹⁶ An argument similar to that used in Lem. 4 makes this explicit.

¹⁷ But $\phi_{\overline{G}_\varepsilon}$ is not a subformula of $(\mu X \cdot \phi)$, which is why we do not use structural induction.

Since $f_{\overline{\mathcal{G}}_\varepsilon} . x \leq f_{\underline{\mathcal{G}}', \overline{\mathcal{G}}_\varepsilon} . x$ we have $x \ll f_{\underline{\mathcal{G}}', \overline{\mathcal{G}}_\varepsilon} . x$ — and we recall from Lem. 8 that $f_{\underline{\mathcal{G}}', \overline{\mathcal{G}}_\varepsilon}$ is ok. Hence $(\sqcup n \cdot f_{\underline{\mathcal{G}}', \overline{\mathcal{G}}_\varepsilon}^n . x) = \mu . f_{\underline{\mathcal{G}}', \overline{\mathcal{G}}_\varepsilon}$, and we can conclude with

$$(\sqcup n \cdot f_{\overline{\mathcal{G}}_\varepsilon}^n . x) \leq (\sqcup n \cdot f_{\underline{\mathcal{G}}', \overline{\mathcal{G}}_\varepsilon}^n . x) = \mu . f_{\underline{\mathcal{G}}', \overline{\mathcal{G}}_\varepsilon} = \mu . f_{\overline{\mathcal{G}}_\varepsilon} .$$

Thus $f_{\overline{\mathcal{G}}_\varepsilon}$ is indeed ok.