



# A Novel Stochastic Game Via the Quantitative $\mu$ -calculus

Annabelle McIver<sup>1,2</sup>

*Dept. Computer Science  
Macquarie University  
Sydney NSW 2109, Australia*

Carroll Morgan<sup>1,3</sup>

*Dept. Comp. Sci. & Eng.  
University of New South Wales  
Sydney NSW 2052, Australia*

---

## Abstract

The quantitative  $\mu$ -calculus  $qM\mu$  extends the applicability of Kozen's standard  $\mu$ -calculus [5] to probabilistic systems. Subsequent to its introduction [9,4] it has been developed by us [6,7,8] and by others [2]. Beyond its natural application to define probabilistic temporal logic [10], there are a number of other areas that benefit from its use.

One application is stochastic two-player games, and the contribution of this paper is to depart from the usual notion of “absolute winning conditions” and to introduce a novel game in which players can “draw”.

The extension is motivated by examples based on economic games: we propose an extension to  $qM\mu$  so that they can be specified; we show that the extension can be expressed via a reduction to the original logic; and, via that reduction, we prove that the players can play optimally in the extended game using memoryless strategies.

*Keywords:* Probabilistic systems, mu-calculus, quantitative logic, stochastic games, intermediate fixed points, draw and stalemate.

---

---

<sup>1</sup> We gratefully acknowledge the support of the Australian Research Council (ARC) via its Discovery Project Grant DP034557.

<sup>2</sup> Email: [anabel@ics.mq.edu.au](mailto:anabel@ics.mq.edu.au)

<sup>3</sup> Email: [carrollm@cse.unsw.edu.au](mailto:carrollm@cse.unsw.edu.au)

# 1 Introduction

Probabilistic systems combine standard features of computer systems and random events, thus a significant number of their properties are quantitative rather than qualitative and cannot be verified using standard methods. *Quantitative program logics* [4,9,14] have been developed to overcome the limitations of ordinary program logic by specifically taking into account probabilistic actions. In particular the *quantitative  $\mu$ -calculus*  $qM\mu$  extends the standard  $\mu$ -calculus of Kozen [5], giving access to probabilities by interpreting terms as real- rather than Boolean-valued functions of the state. As in the standard  $\mu$ -calculus,  $qM\mu$  comprises a modelling language with a simple denotational semantics, as well as an operational interpretation as a two-player game extending Stirling [12] and establishing a formal link to *stochastic parity games* [1] and their associated algorithmic methods of verification.

However not all games fit naturally within this framework, and in this paper we investigate one novel game which arises in the specification of some (not necessarily probabilistic) systems, and show how the quantitative  $\mu$ -calculus nevertheless can successfully accommodate it.

As an example of the type of games we study, we consider the situation in which two players called *Max* and *Min* are given a pile of twenty \$1 coins to share between each other. They agree to execute the following protocol to ensure a fair division: first *Min* chooses an amount  $s$  to give to *Max* who either accepts or rejects it. In the case that he accepts, *Min* then receives  $\$(20-s)$ ; otherwise she is forced to choose again. Essentially there are only two outcomes of this game — either play ends after some finite time with *Max* and *Min* having divided the coins to their mutual satisfaction; or they play indefinitely, never able to agree.

In standard (and quantitative)  $\mu$ -calculus, infinite executions are dealt with via fixed points — essentially the greatest fixed point is interpreted as an *absolute* win for *Max* and the least fixed point the opposite, *i.e.* an absolute win for *Min*. In the above protocol however, it seems clear that the infinite execution should be judged as some kind of “draw” or at least a “stalemate”, rather than a decisive win for either player, an outcome which cannot be modelled by either a single least- or greatest fixed point. The solution we adopt leads us to propose a new kind of fixed point which we identify with intermediate winning conditions. In particular our contributions can be summarised as follows.

- (i) A simple extension to the quantitative  $\mu$ -calculus allowing games exhibiting intermediate “draws” to be specified directly (Sec. 3 and Sec. 4);
- (ii) A demonstration that such games are indeed well defined in the traditional sense of game theory, by which we mean that the players can assess

the effect of their strategies relative to those of the other player and that the two players can each follow optimal strategies (Sec. 4);

- (iii) A detailed case study inspired by an economic application, and a method for analysis based on the above theoretical results (Sec. 5).

Crucial to the approach will be the use of  $qM\mu$ , and in Sec. 2 we review its interpretation over probabilistic systems.

Throughout, the following notational conventions apply. An infix ‘.’ is used for function application. We write  $S$  for a (fixed) underlying state space, and  $\overline{S}$  for the set of *discrete probability distributions* over  $S$ , where a discrete probability distribution is a function from  $S$  to the interval  $[0, 1]$  which is normalised to 1; thus  $\overline{S} \hat{=} \{\Delta: S \rightarrow [0, 1] \mid \sum_{s \in S} \Delta.s = 1\}$ . The set of functions from  $S$  to the real interval  $[0, 1]$  is denoted by  $\mathcal{E}S$ , and called the *expectations* over  $S$ . Real-valued functions over  $S$  (e.g. expectations) are ordered by lifting the pointwise the order  $\leq$  on the reals; functions  $\max$  and  $\min$  are similarly lifted; and we write  $\underline{x}$  for the constant function returning some real  $x$  for all states, where usually we will have  $0 \leq x \leq 1$ . If  $\Delta$  is a probability distribution over  $S$  and  $A$  is a measurable function on  $S$  then  $\int_{\Delta} A$  denotes the expected value of  $A$  with respect to  $\Delta$ . When  $\Delta$  is in  $\overline{S}$  and  $A$  is a bounded real-valued function over  $S$ , this reduces to  $\sum_{s \in S} \Delta.s \times A.s$ .

## 2 Quantitative $qM\mu$ -calculus and games

In this section we summarise the details of the  $\mu$ -calculus, beginning with the definition of the language, and then reviewing how formulae can be interpreted in two equivalent ways over probabilistic transition systems. Formulae  $\phi$  in the logic (in positive form) are constructed as follows:

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

In the interpretations of the formulae, the following meanings will be given.

- Variables  $X$  are used for binding fixed points.
- Terms  $\mathbf{A}$  stand for fixed functions (normally) in  $\mathcal{E}S$ .
- Terms  $\mathbf{k}$  represent labelled probabilistic transitions (described below).
- Terms  $\mathbf{G}$  describe Boolean functions of  $S$ , i.e. (“if”)  $\triangleleft \mathbf{G} \triangleright$  (“else”).
- $(\mu X \cdot \phi)$  and  $(\nu X \cdot \phi)$  are extremal fixed points, binding any free  $X$ ’s in  $\phi$ .

We avoid use of the usual modalities  $\langle \cdot \rangle$  and  $[\cdot]$  forming respectively angelic (existential-) and demonic (universal-) choices, as they can be expressed equivalently (in the assumptions of our framework) by a combination of  $\{\mathbf{k}\}\phi$  and

$\sqcup$  and  $\sqcap$ .<sup>4</sup>

Next we show how to interpret the above language over a *probabilistic transition system*. Such a transition system is modelled by functions from initial states in  $S$  to final distributions in  $\overline{S}$ , where the distributions model probabilistic features present in the system. In this paper we shall use labels to distinguish the various probability distributions within a given result set — thus our abstract computational model is (strictly speaking) *labelled probabilistic transition systems*.<sup>5</sup> Let  $L$  be a (finite) index set of labels; we write  $\mathcal{R}.L.S$  for a labelled probabilistic transition system over  $S$ , so that it has the type  $L \rightarrow S \rightarrow \overline{S}$ . Thus given a label and an initial state, the result is a single output distribution in  $\overline{S}$ , which we call a *transition probability distribution*.

In order to interpret a formula over a given (labelled) probabilistic transition system  $r \in \mathcal{R}.L.S$ , we use the standard technique of *valuations* from denotational semantics [13] which works roughly as follows. Given a formula  $\phi$ , a valuation  $\mathcal{V}$  does four things: (i) it maps each  $\mathbf{A}$  in  $\phi$  to a fixed expectation in  $\mathcal{E}S$ ; (ii) it maps each  $\mathbf{k}$  to a fixed label, and thus  $\{\mathbf{k}\}$  to the corresponding fixed probabilistic transition determined by  $r$ ; (iii) it maps each  $\mathbf{G}$  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.)

Formulae can be interpreted in two equivalent ways [6]: denotationally, extending Kozen [5], or operationally as a game, extending Stirling [12]. We now present each in turn.

**The denotational interpretation** gives the meaning of a formula  $\phi$  as a function in  $\mathcal{E}S$ , generalising Kozen’s interpretation as a Boolean function of  $S$  [5]. Let  $\phi$  be a formula and  $\mathcal{V}$  a valuation. We write  $\|\phi\|_{\mathcal{V}}$  for its meaning determined by the rules given in Fig. 1.

**The operational interpretation** of a formula  $\phi$  is in terms of generalisation of Stirling’s turn-based games [12], which we call *Stochastic Stirling Games* (*SSG*’s). The game is between two players we call *Max* and *Min*; *Max*’s objective is to maximise a certain “payoff” (defined below) and *Min*’s is to minimise it. Play progresses through a sequence of *game positions*, each

<sup>4</sup> The modality  $\langle K \rangle \phi$  for example is equivalent to the angelic choice over all terms of the form  $\{\mathbf{k}\}\phi$ , for  $\mathbf{k}$  in the (finite) subset  $K$ .

<sup>5</sup> In this paper, however, labels play no role except as convenient “markers” to distinguish different distributions; they are not, for example, used to express path properties of the computational sequences.

- (i)  $\|A\|_{\mathcal{V}} \hat{=} \mathcal{V}.A$  .
- (ii)  $\|\{k\}\phi\|_{\mathcal{V}.s} \hat{=} \int_{\mathcal{V}.k.s} \|\phi\|_{\mathcal{V}}$  .
- (iii)  $\|\phi' \sqcap \phi''\|_{\mathcal{V}.s} \hat{=} \|\phi'\|_{\mathcal{V}.s} \min \|\phi''\|_{\mathcal{V}.s}$  ;  
 and  $\|\phi' \sqcup \phi''\|_{\mathcal{V}.s} \hat{=} \|\phi'\|_{\mathcal{V}.s} \max \|\phi''\|_{\mathcal{V}.s}$  .
- (iv)  $\|\phi' \triangleleft G \triangleright \phi''\|_{\mathcal{V}.s} \hat{=} \|\phi'\|_{\mathcal{V}.s} \underline{\text{if}} (\mathcal{V}.G.s) \underline{\text{else}} \|\phi''\|_{\mathcal{V}.s}$  .
- (v)  $\|(\mu X \cdot \phi)\|_{\mathcal{V}} \hat{=} (\text{lfp } \varepsilon \cdot \|\phi\|_{\mathcal{V}_{[X \mapsto \varepsilon]}})$  , where by  $(\text{lfp } \varepsilon \cdot \text{exp})$  we mean the least fixed-point of the function  $(\lambda \varepsilon \cdot \text{exp})$  in  $\mathcal{ES} \rightarrow \mathcal{ES}$ .
- (vi)  $\|(\nu X \cdot \phi)\|_{\mathcal{V}} \hat{=} (\text{gfp } \varepsilon \cdot \|\phi\|_{\mathcal{V}_{[X \mapsto \varepsilon]}})$  .

Note that in the valuation  $\mathcal{V}_{[X \mapsto \varepsilon]}$ , the variable  $X$  is re-mapped to the expectation  $\varepsilon$ , and the fixed points exist because the least- and greatest expectations in  $\mathcal{ES}$  are  $\underline{0}$  and  $\underline{1}$ .

Fig. 1. Kozen-style denotational semantics for  $qM\mu$

of which is either a pair  $(\phi, s)$  where  $\phi$  is a formula and  $s$  is a state in  $S$ , or is a single  $(y)$  for some real-valued payoff  $y$  in  $[0, 1]$ . Following Stirling, we will use “colours” to handle repeated returns 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  $(y)$  at the end. The initial formula  $\phi_0$  is the given  $\phi$ , and  $s_0$  is an *initial* state in  $S$ . A move from position  $(\phi_i, s_i)$  to  $(\phi_{i+1}, s_{i+1})$  or to  $(y)$  is specified by the rules of Fig. 2.

Imagine that the two players play according to the rules in Fig. 2, given a formula  $\phi$  and an initial state  $s_0$ , and suppose first that there are no probabilistic transitions. In this simple case the result would be a single game path recording the actual sequence of game positions observed during a play, where the players have decided “on the fly” how to resolve their choices as they go along. Alternatively, and to achieve the same effect, they could formulate a strategy *beforehand* to resolve their choices depending on how the game has progressed so far. We call such strategies *pre-determined*, and we model them as functions  $\sigma : \text{paths} \rightarrow \{\text{true}, \text{false}\}$ , *i.e.* from finite game paths to Booleans. We use the convention that **true** corresponds to “take the left branch at a  $\sqcap$  (or  $\sqcup$ )”, and **false** to “take the right branch at a  $\sqcap$  (or  $\sqcup$ )”. The idea is that at each relevant game position the player “looks up” his strategy after a particular (finite) play which then determines whether to continue play by following the left- or the right branch. As the strategies take paths as arguments, they model *unlimited memory*. We write  $\underline{\sigma}$  and  $\overline{\sigma}$  for pre-determined strategies as would be used by *Min* and *Max* respectively.

Now suppose more generally that the transitions can be probabilistic; this means that some of the moves are determined by chance. Nevertheless the players can still play according to their respective strategies, but this time the

If the current game position is  $(\phi, s)$ , then play proceeds as follows:

1. If  $\phi$  is **A** then the game terminates in position  $(y)$  where  $y = \mathcal{V}.A.s$ .
2. if  $\phi$  is  $\{k\}\phi$  then the distribution  $\mathcal{V}.k.s$  is used to choose a next state  $s'$  in  $S$ : the next game position is  $(\phi, s')$  (with probability  $\mathcal{V}.k.s.s'$ ).
3. If  $\phi$  is  $\phi' \sqcap \phi''$  (resp.  $\phi' \sqcup \phi''$ ) then *Min* (resp. *Max*) chooses one of the minjuncts (maxjuncts): the next game position is  $(\phi, s)$ , where  $\phi$  is the chosen 'junct  $\phi'$  or  $\phi''$ .
4. If  $\phi$  is  $\phi' \triangleleft G \triangleright \phi''$ , the next game position is  $(\phi', s)$  if  $\mathcal{V}.G.s$  holds, and otherwise it is  $(\phi'', s)$ .
5. If  $\phi$  is  $(\mu X \cdot \phi)$  (resp.  $(\nu X \cdot \phi)$ ) then a fresh colour **C** is chosen and is bound to the formula  $\phi_{[X \mapsto C]}$ , in which  $X$  has been replaced by **C**, for later use; the next game position is  $(C, s)$ . (This use of colours makes easy determination of which recursion operator actually “caused” an infinite path — see below.)
6. If  $\phi$  is a colour **C**, then the next game position is  $(\Phi, s)$ , where  $\Phi$  is the formula bound previously to **C**.

The game begins with a closed formula. Infinite plays *always* result in there being exactly one colour **C** that occurs infinitely often [12,8].

The outcome of the game is determined by a *payoff* function called *Val*, which is defined as follows. Note that it is both insensitive to the colours and length on finite paths.

7. The path  $\pi$  is finite, terminating in a game state  $(y)$ ; in this case the value  $Val.\pi$  is  $y$ .
8. The path  $\pi$  is infinite and there is (exactly one) colour **C** appearing infinitely often that was generated by a greatest (resp. least) fixed-point  $\nu$  (resp.  $\mu$ ); in this case  $Val.\pi$  is 1 (resp. 0).

Fig. 2. Rules for the Stochastic Stirling Game.

result of their doing so is not a single possible game path, but a set of paths, each one labelled (at least for finite paths) with the probability of occurrence in a play of the game. More generally the structure represents a *probability distribution over game paths*. For example let the formula be

$$(1) \quad \psi \hat{=} \{s := s_0 \text{ }_{1/2} \oplus \text{ } s := s_1\}(\underline{1} \sqcap (\mu X \cdot X)) ,$$

where we have instantiated a particular transition (rather than writing  $k$ ), allowing it to stand for itself, and we use  $s_0$  and  $s_1$  for states in  $S$ . We are also using  $P \text{ }_p \oplus Q$  to denote a probabilistic transition in which  $P$  is selected with probability  $p$  and  $Q$  with probability  $1-p$ . Now let *Min*'s strategy be “if the state  $s$  is  $s_0$  take the left branch, otherwise take the right branch”; we use

C for the colour binding the least fixed point. The two paths generated when *Min* plays according to her strategy are

$$\begin{aligned} \pi_0 &\hat{=} ((\psi, s_0), (\underline{1} \sqcap (\mu X \cdot X), s_0), (1)) , \text{ and} \\ \pi_1 &\hat{=} ((\psi, s_0), (\underline{1} \sqcap (\mu X \cdot X), s_1), ((\mu X \cdot X), s_1), (C, s_1), (C, s_1), \dots) , \end{aligned}$$

both occurring with probability 1/2. In fact the probability distribution over game paths is defined by the well-known measure [3] generated by the particular probabilistic transitions of the underlying transition system. We call the so-generated distribution a *path distribution*. Clearly a different path distribution will be generated if *Min* follows a different strategy. More generally, given strategy sequences  $\underline{\sigma}$  and  $\bar{\sigma}$  we write  $\llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}} . s_0$  for the resulting path distribution over game paths when *Min* and *Max* play according to  $\underline{\sigma}$  and  $\bar{\sigma}$  respectively starting from  $s_0$ .

Next, we consider the payoff function *Val*. According to the rules in Fig. 2 we see in our example at (1) that  $Val.\pi_0 = 1$ , since it is a finite path with final constant term 1, but that  $Val.\pi_1 = 0$ , since the infinitely-occurring colour C was generated by a least fixed point. Thus the expected payoff with respect to the distribution over game paths with *Min* playing the above strategy is  $1/2 \times Val.\pi_0 + 1/2 \times Val.\pi_1 = 1/2$ . More generally we write  $\int \llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}} . s_0$  for the expected payoff with respect to the path distribution generated when *Min* and *Max* follow strategy sequences  $\underline{\sigma}$  and  $\bar{\sigma}$  from initial state  $s_0$ . It is well defined since *Val* is integrable over the  $\sigma$ -algebra of game paths [14].

We say that the game *as a whole* is well defined (in the sense that each player can play rationally, *i.e.* can assess a particular strategy relative to all the other player’s strategies), if the *minimax* over all possible strategy sequences of the expected payoff is equal to the *maximin*. That is, for all  $s$  in  $S$  we must have

$$(2) \quad \sqcap_{\underline{\sigma}} \sqcup_{\bar{\sigma}} \int \llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}} . s \quad = \quad \sqcup_{\bar{\sigma}} \sqcap_{\underline{\sigma}} \int \llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}} . s .$$

We call this the *value* of the game, and often refer to it as  $V^*$ . A strategy  $\underline{\sigma}^*$  is *optimal* for *Min* if  $\int \llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}^*} . s \leq V^*$  for all *Max* strategies  $\bar{\sigma}$ ; similarly a strategy  $\bar{\sigma}^*$  is *optimal* for *Max* if  $\int \llbracket \phi \rrbracket_{\mathcal{V}}^{\underline{\sigma}^*, \bar{\sigma}^*} . s \geq V^*$  for all *Min* strategies  $\underline{\sigma}$ .

The following results summarise the relationship between *SSG*’s and the denotational interpretation of  $qM\mu$ ; they are both proved elsewhere [6,7,8].

**Theorem 2.1** *In finite state spaces, a Stochastic Stirling Game given by formula  $\phi$  is well defined and its value  $V^*$  is equal to its denotation  $\llbracket \phi \rrbracket_{\mathcal{V}}$ .*

Thus, for  $\psi$  defined at (1), the associated *SSG* has value 0, since that is its denotation according to the rules set out in Fig. 1.

Amongst player strategies, the memoryless ones are particularly significant, where a strategy is said to be *memoryless* if it is independent of the history, only depending on the current game position.

**Theorem 2.2** *In finite state spaces  $S$ , each player playing an *SSG* has a memoryless optimal strategy.*

We usually write  $\underline{\sigma}^*$  and  $\overline{\sigma}^*$  for the optimal strategies.

One consequence of Thm. 2.2 is that, given a valuation  $\mathcal{V}$ , each specific  $\sqcup$  and  $\sqcap$  in a formula  $\phi$  may be replaced by a specific Boolean choice, where the choice depends only on the underlying state space  $S$ , so that all “on-line” choices made by either player can be removed. When players play memoryless strategies the resulting probabilistic structure is essentially a Markov chain [3].

For the game given by  $\psi$  at (1) *Min*’s optimal strategy is “always take the right branch”, which is equivalent to replacing the  $\sqcap$  in  $\psi$  with  $(\cdot \triangleleft \text{false} \triangleright \cdot)$  to yield a formula  $\psi' \hat{=} \{x := 0\}_{1/2} \oplus \{x := 1\}(\underline{1} \triangleleft \text{false} \triangleright (\mu X \cdot X))$ . Observe that the formerly present choices of *Min* have been effectively removed, forcing her to follow the instantiated memoryless strategy — Thm. 2.2 guarantees that the value of the new  $\psi'$ -game is the same as the value of the original  $\psi$ -game. As we see, the instantiated Boolean choice always selects the infinite iteration, with a payoff of 0 for both paths.

More generally if  $\underline{\sigma}$  ( $\overline{\sigma}$ ) represents a memoryless strategy for *Min* (*Max*) then we write  $\phi^{\underline{\sigma}}$  ( $\phi^{\overline{\sigma}}$ ) for the formula with each syntactic occurrence of  $\sqcap$  ( $\sqcup$ ) replaced by the Boolean choices determined by  $\underline{\sigma}$  ( $\overline{\sigma}$ ) as illustrated above.

### 3 Generalising the payoff for infinite play

In this section we introduce our generalisation of *SSG*’s — we call them *Asymptotic SSG*’s. Roughly speaking, the idea is to extend the definition of *Val* so that along specified infinite paths the payoff is not restricted to be 0 or 1, but can take some other (constant) value  $x$  lying strictly in between.

To illustrate we return to the game between *Max* and *Min* described in the introduction. First we encode a single move of the protocol in *qMμ* using the rules in Fig. 2 as a guide. Let *Min*’s choice be given by the transition *Choose*  $\hat{=} s \underline{::} \{0 \dots 20\}$ , where we are using the shorthand “ $\underline{::}$ ” to indicate a minimising choice from the set concerned. Next *Max* has the choice of accepting  $s$ , or rejecting it and forcing *Min* to choose again. We define the real-valued function **Accept** to return the value of the choice offered in the

current state  $s$ , so that if  $s = 3$  then in fact  $\text{Accept}.s = 3$  as well. (That is, because the state itself is real-valued, the constant expectation is the identity in this case.) Thus we can model  $Max$ 's move with the sub-formula  $\text{Accept} \sqcup X$ , where  $X$  is a variable to be bound to a fixed point to model the repetition in the case of rejection. In all, that suggests a formula  $(\tau X \cdot \{\text{Choose}\}(\text{Accept} \sqcup X))$ , formally describing the desired protocol for dividing the money — but where we have introduced an undefined-for-now fixed point binding  $\tau$  for  $X$ , as we try to decide which one to use.

The first thing to notice about our proposed formula is that the payoff is not constrained to lie between  $[0, 1]$ , which was one of the assumptions set out in the rules of the game in Fig. 2: for example  $Min$  might choose  $s = 20$  and  $Max$  might accept. However this is not a significant issue provided that all the payoffs lie within a bounded and closed subset of the reals. In such cases we can always transform the game using an *affine transformation*,<sup>6</sup> which can shift and scale the payoff functions without changing the underlying probabilistic transitions. In this case all we need do is scale  $\text{Accept}$  by  $1/20$ . The next lemma sets out the details in general.

**Lemma 3.1** *Consider a formula  $\phi$ , and a valuation which maps constant expectations  $A$  in  $\phi$  to some interval  $[a, b]$  (with  $a < b$ ). If  $Val$  awards a payoff of  $a$  to infinite paths won by  $Min$ , and  $b$  to infinite paths won by  $Max$ , then the value of the game is  $\alpha^{-1} \cdot (\|\phi\|_{(\alpha, \mathcal{V})})$ , where  $\alpha$  is the affine transformation mapping  $[a, b]$  onto  $[0, 1]$  and the term  $\alpha \cdot \mathcal{V}$  is defined so that only the constant terms specified in  $\mathcal{V}$  are transformed; the transition probabilities are unchanged.*

Even with the transformed game  $(\tau X \cdot \{\text{Choose}\}((\text{Accept}/20) \sqcup X))$ , we still have the problem of deciding which fixed point to use. A *greatest* fixed point would award the maximum 1 for an infinite play, corresponding to  $Max$  taking all 20 coins; on the other hand a *least* fixed point would award to an infinite play a value of 0, allowing the bias to favour  $Min$  so that she wins everything.

Our intention is that the stalemate situation should correspond to an equal split — thus effectively  $Min$  is forced to divide the pile of coins into two equal parts, which  $Max$  must accept. That corresponds to a game which allows infinite plays to be awarded a constant value of  $1/2$ ; put another way, this encourages us to define a new kind of fixed point where infinite iteration is awarded  $1/2$  rather than 0 or 1. Thus we firm up our new fixed-point notation,

<sup>6</sup> An affine transformation  $t$  is a combination of scaling by a fixed factor and a shift by a constant amount, and can be defined  $t.e \hat{=} \gamma \times e + \delta$ . We restrict to affine transformations where the scaling factor  $\gamma$  is non-negative.

explicitly annotating the  $\tau$  with the value it is to award to infinite plays, *i.e.*

$$(\tau_{1/2} X \cdot \{Choose\}((Accept/20) \sqcup X)) ,$$

where in general a  $\tau_x$  indicates that its infinite plays are awarded  $x$ .

In the next section we turn to the main contribution of this paper: we provide the necessary theory to show that games with such intermediate payoffs for infinite paths are indeed well defined, and that the players have optimal strategies for them. Crucial to our argument is to show that they have an (unusual) denotational interpretation in  $qM\mu$ , for that then lets us appeal immediately to Thms. 2.1 and 2.2.

### 4 Extended $qM\mu$ and Asymptotic $SSG$ 's

With Lem. 3.1, we can without loss of generality work within the interval  $[0, 1]$ . We begin by defining a new language constructor  $\tau_x$  which provides the “intermediate” type of loop back introduced informally above.

**Definition 4.1** We extend the quantitative  $\mu$ -calculus by adding the new constructor  $(\tau_x X \cdot \phi)$  to the earlier definition in Sec. 2.<sup>7</sup> Here  $x$  is any real satisfying  $0 \leq x \leq 1$ .

Next, as in Fig. 2, we can define similarly a game based on unfolding any expression in the extended quantitative  $\mu$ -calculus — indeed the rules for plays are unchanged, with the only substantial modification being the winning condition relative to the new constructor.

**Definition 4.2** Given an expression in the extended  $\mu$ -calculus, the associated *Asymptotic SSG* is played as set out in Fig. 2, with the following modifications.

- 5'. If  $\phi$  is  $(\tau_x X \cdot \phi)$  then a fresh colour  $C$  is chosen and is bound to the formula  $\phi_{[X \mapsto C]}$  for later use; the next game position is  $(C, s)$ .
- 9. If a colour belonging to a  $\tau_x$  occurs infinitely often along a path, the payoff is  $x$ .

Unfortunately for our putative Def. 4.2 there remains the question of well-definedness when we consider the minimax value of such a game — recall that this relies on an equality of the form (2). Given strategy sequences, the expected payoff  $\int_{[\phi]_{\nu}^{\sigma, \sigma}.s} Val$  is indeed well defined [6] — however it is not immediate that the minimax of expected payoffs is equal to the maximin. Below we show nevertheless that they are equal, and give methods for computing the value.

---

<sup>7</sup> At the same time we could if we wished dispense with the original least- and greatest fixed-point operators  $\mu, \nu$ , replacing them by  $\tau_0, \tau_1$  respectively.

Our approach is to reduce the asymptotic game to a standard *SSG*, by expressing it as a  $qM\mu$  formula, and then appealing to Thm. 2.1. Throughout we work with a fixed valuation  $\mathcal{V}$ . For simplicity we shall only consider formulae of the form  $\Phi_A \hat{=} (\tau_x X \cdot \Phi)$ , where  $\Phi$  is fixed-point free and contains no variables other than  $X$ . We will refer to the associated game as the *asymptotic game*; the next definition gives a standard  $qM\mu$  formula which we shall show is equivalent to it.

**Definition 4.3** Given a  $qM\mu$  formula  $\Phi$  with a single free variable  $X$  (but possibly occurring in several positions), so that  $(\tau_x X \cdot \Phi)$  is in the extended quantitative  $\mu$ -calculus, we define the *x-gadget* game to be the *SSG* corresponding to the  $qM\mu$  formula

$$\Phi_G \hat{=} (\nu Y \cdot (\mu Z \cdot \Phi(Z \sqcup (Y \sqcap \underline{x})))) ,$$

where in general by  $\Phi(expr)$  we mean  $\Phi_{[X \mapsto expr]}$ .<sup>8</sup>

Although Def. 4.3 seems quite odd, the interleaving of the fixed points allows the players essentially to “simulate” the asymptotic game; Def. 4.3 is also a fixed point of (the denotation of)  $\Phi(\cdot)$ , although in general neither the greatest nor the least. Intuitively we have specified a “get out and take  $x$ ” clause, ensuring that it is controlled by neither player independently — the extra decision to be resolved by *Max* ensures that the game does not end prematurely, resulting in a payoff that is *too small*. Similarly *Min* can ensure that the payoff is not *too large*. The only remaining case is for *Min* to terminate the game for an immediate payoff of  $x$  — but as we shall see that happens exactly when the original asymptotic game would have resulted in an infinite  $\tau_x$ -path when both players play optimally. Next we consider how to formalise the above argument. We proceed in two stages, summarised as follows.

- First we define a new game via formula  $\Phi_H = (\tau_x Y \cdot (\mu Z \cdot \Phi(Z \sqcup Y)))$ ; it corresponds directly to the asymptotic game,  $\Phi_A$ . We show that the two games generate isomorphic path distributions, and therefore that their expected payoffs correspond.
- Second we show that there is similarly a direct correspondence between  $\Phi_H$  and the *x-gadget* game,  $\Phi_G$ , defined above.

Turning now to the details, the following two lemmas deal with the first

<sup>8</sup> We note that this technique of reducibility — forcing *Min* and *Max* to compete via the interleaved fixed points and extra decision points — is reminiscent of the “gadgets” invented by Henzinger *et. al.* for converting a stochastic parity game to a standard (non-stochastic) parity game [1], with corresponding, though not equivalent properties. Here the correspondence turns out to be exact.

stage, and set up the correspondence between  $\Phi_A$  and  $\Phi_H$ .

**Lemma 4.4** *Given any strategy sequences  $\underline{\sigma}$  and  $\bar{\sigma}$  for  $\Phi_A$ , there exist strategy sequences  $\underline{\sigma}'$  and  $\bar{\sigma}'$  for  $\Phi_H$  such that the resulting path distributions correspond.*

**Proof (Sketch)** Comparing the syntactic construction given in Fig. 2 of  $\Phi_A$  and  $\Phi_H$  we observe only two differences. The first is that the extra fixed point in  $\Phi_H$  produces extra “colour” positions, and the second is that the extra  $\sqcup$  at  $Z \sqcup Y$ , also in  $\Phi_H$ , leads to an extra choice for *Max*. Nevertheless most of the game moves are exactly the same: that is because all moves defined by  $\Phi$  — which make up the bulk of the activity — are identical in the two formulae. In particular the finite paths ending in a constant term must be determined by navigating through  $\Phi$ . Thus we can define  $\underline{\sigma}'$  and  $\bar{\sigma}'$  by forcing the players of  $\Phi_H$  to copy the decisions encoded in  $\underline{\sigma}$  and  $\bar{\sigma}$  for unwrapping  $\Phi$ ; the resolution of the extra  $\sqcup$  at  $Z \sqcup Y$  can be arbitrary. The result is that the proportions of finite and infinite paths must be the same in the path distributions for the two games (since the option to repeat is determined within  $\Phi$ ); moreover each finite path in  $\llbracket \Phi_A \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}}.s$  corresponds to a finite path in  $\llbracket \Phi_H \rrbracket_{\mathcal{V}}^{\underline{\sigma}', \bar{\sigma}'}.s$ , terminating with the same real. The result follows.  $\square$

The next lemma shows that the *expected payoffs* correspond with respect to the payoff functions.

**Lemma 4.5** *Let  $\underline{\sigma}, \bar{\sigma}$  be strategy sequences for  $\Phi_A$ , and let  $\underline{\sigma}', \bar{\sigma}'$  be derived strategies for  $\Phi_H$  as in Lem. 4.4. The following inequality holds:*

$$(3) \quad \int_{\llbracket \Phi_A \rrbracket_{\mathcal{V}}^{\underline{\sigma}, \bar{\sigma}}.s} Val \geq \int_{\llbracket \Phi_H \rrbracket_{\mathcal{V}}^{\underline{\sigma}', \bar{\sigma}'}.s} Val .$$

**Proof (Sketch)** By Lem. 4.4 the two path distributions correspond, thus the inequality at (3) will follow provided that the payoff function awards values in the  $\Phi_A$  case are at least as great as those it awards in the  $\Phi_H$  case. But this follows immediately, since by Lem. 4.4 (again) the finite paths correspond and from Fig. 2 the payoffs on those paths are the same. Considering now the payoffs for infinite paths, we see that  $\Phi_A$ , with only one fixed point  $\tau_x$ , generates infinite paths containing a single colour, thus it must be awarded a payoff of  $x$  by Def. 4.2. On the other hand  $\Phi_H$ , with two fixed points, generates infinite paths in which the infinite colour was generated either by a  $\mu$  or a  $\tau_x$ . In the former case, *Val* awards 0 (recall Fig. 2) and in the latter case it awards  $x$  (recall Def. 4.2). Either way the payoff is no more than  $x$  and the result now follows by monotonicity.  $\square$

Now we have accomplished the first stage of the argument, our next task

is to show how  $\Phi_H$  and the gadget-game  $\Phi_G$  correspond.

**Lemma 4.6** *Let  $C$  be the colour bound to the outermost fixed point in  $\Phi_H$ . Given strategy sequences  $\underline{\sigma}$  and  $\bar{\sigma}$  for  $\Phi_H$  there exists a strategy sequence  $\bar{\sigma}'$ , and for each  $n \geq 0$  a strategy sequence  $\underline{\sigma}'_n$  such that*

$$(4) \quad \int_{\llbracket \Phi_H \rrbracket_{\bar{v}}^{\underline{\sigma}, \bar{\sigma}} .s} Val_n = \int_{\llbracket \Phi_G \rrbracket_{\bar{v}}^{\underline{\sigma}'_n, \bar{\sigma}' } .s} Val ,$$

where  $Val_n$  is defined to be  $Val$  on paths which have no more than  $n$  occurrences of the colour  $C$ , and  $x$  otherwise.

**Proof (Sketch)** As in Lem. 4.4 we see that the syntactic rules from Fig. 2 applied to  $\Phi_H$  and  $\Phi_G$  are almost the same. The only differences this time are that the outer-most fixed point is  $\tau_x$  in  $\Phi_H$  and  $\nu$  in  $\Phi_G$ , and that there is an extra  $\sqcap \underline{x}$  in  $\Phi_G$ . The effect of the latter difference is that  $Min$  playing  $\Phi_G$  has the additional option to terminate the game in a finite play for an immediate payoff of  $x$ . For all other cases (when she does not decide to terminate there), the players playing  $\Phi_G$  can simply copy all the moves made by the players playing  $\Phi_H$ . Thus, as before, we define  $\bar{\sigma}'$  to be the strategy which copies all the corresponding decisions encoded in  $\bar{\sigma}$ .

To define  $\underline{\sigma}'_n$  we must take  $Val_n$  into account. Consider any game path generated by  $\Phi_H$  which has more than  $n$  occurrences of a colour  $C$  (for arbitrary states  $s$ ); by definition  $Val_n$  awards each such path a value of  $x$ . To obtain the equivalence at (4) above we need to show that the corresponding paths in  $\Phi_G$  can also be awarded a payoff of  $x$ . We do that by exploiting the extra  $\sqcap \underline{x}$  in  $\Phi_G$  to terminate the  $\Phi_G$ -game early for those paths. Note first that since outermost fixed points generate a single colour, there is a direct correspondence between the number of  $C$ -colour positions in  $\Phi_H$  and the number of  $C'$ -colour positions in  $\Phi_G$ , where  $C'$  is the colour bound to the outer-most fixed point in  $\Phi_G$ . Define  $\underline{\sigma}'_n$  to behave like  $\underline{\sigma}$  for all decisions *not* involving the extra  $\sqcap$ , and to select the left branch at the extra  $\sqcap$  (*i.e.* not to terminate early) if there are fewer than  $n$   $C'$ -colour positions, and to select the right branch otherwise (*i.e.* terminate the game early for an immediate payoff of  $x$ ). The result now follows.  $\square$

Having completed the second stage, we can finally prove our desired correspondence.

**Lemma 4.7** *Let  $\bar{\sigma}^*$  and  $\underline{\sigma}^*$  be the optimal strategies for the  $x$ -gadget game  $\Phi_G$ , and let  $\underline{\sigma}$  and  $\bar{\sigma}$  be arbitrary strategy sequences for the asymptotic game*

$\Phi_A$ ; further let  $V^*$  be  $\|\Phi_G\|_{\mathcal{V}}$  as before. Then the following inequalities hold:

$$\int_{[\Phi_A^*]_{\mathcal{V}}^{\underline{\sigma}.s}} Val \geq V^* \geq \int_{[\Phi_A^*]_{\mathcal{V}}^{\bar{\sigma}.s}} Val$$

where we have used directly  $\underline{\sigma}^*$  and  $\bar{\sigma}^*$  as memoryless strategies for the  $\Phi_A$  game (see explanation below).

**Proof.**

We give the proof only of the first inequality — a dual argument works for the second.

We write  $\Phi_A^*, \Phi_H^*$  and  $\Phi_G^*$  for the formulae with all occurrences of  $\sqcup$  replaced by the (relevant) Boolean-choice determined by the memoryless  $\bar{\sigma}^*$  (we can do that since the  $\sqcup$ 's correspond syntactically in  $\Phi$ , for example, and the replacement simply applies directly the “copying” strategy for *Max* defined at Lem. 4.4 and Lem. 4.6). The result is implied by the following relationships:

$$\int_{[\Phi_A^*]_{\mathcal{V}}^{\underline{\sigma}.s}} Val \geq \int_{[\Phi_H^*]_{\mathcal{V}}^{\underline{\sigma}'.s}} Val = \lim_{n \rightarrow \infty} \int_{[\Phi_H^*]_{\mathcal{V}}^{\underline{\sigma}'_n.s}} Val_n = \lim_{n \rightarrow \infty} \int_{[\Phi_G^*]_{\mathcal{V}}^{\underline{\sigma}''_n.s}} Val \geq V^* .$$

The first inequality follows from Lem. 4.5; the next equality follows from a standard continuity result for integration (specialised to game trees [6]); the next follows from Lem. 4.6, and the last follows since  $\bar{\sigma}^*$  is optimal for *Max* playing  $\Phi_G$ , by definition. □

We can now prove our main theorem, generalising Thm. 2.1 and Thm. 2.2 for games in which infinite plays result in a payoff of  $x$ .

**Theorem 4.8** *The asymptotic game has the same value as the  $x$ -gadget game. That is, the interpretation of  $(\tau_x X \cdot \Phi)$  as determined by Fig. 2 and Def. 4.2 is equal to the interpretation of  $(\nu Y \cdot \mu X \cdot \Phi(X \sqcup (Y \sqcap \underline{x})))$  as determined by Fig. 2 alone.*

*Moreover the players in the asymptotic game have optimal memoryless strategies determined by those for the  $x$ -gadget game.*

**Proof.** Let  $V^*$ ,  $\underline{\sigma}^*$  and  $\bar{\sigma}^*$  be the value, and optimal strategies for the  $x$ -gadget game — they exist by Thms. 2.1 and 2.2. We have the following sequence of inequalities, by appealing first to Lem. 4.7 (recalling that  $\underline{\sigma}$  is arbitrary) and then monotonicity:

$$V^* \leq \sqcap_{\underline{\sigma}} \int_{[\Phi_A^*]_{\mathcal{V}}^{\underline{\sigma}.s}} Val \leq \sqcup_{\bar{\sigma}} \sqcap_{\underline{\sigma}} \int_{[\Phi_A]_{\mathcal{V}}^{\underline{\sigma}.\bar{\sigma}.s}} Val .$$

By similar arguments we also have

$$\sqcap_{\underline{\sigma}} \sqcup_{\overline{\sigma}} \int_{\llbracket \Phi^{\sigma^*} \rrbracket_{\overline{v}.s}} Val \leq V^*$$

which, put together with the earlier inequalities, shows that the maximin dominates the minimax of the asymptotic game. But trivially a minimax dominates a maximin — and so the value  $V^*$  of  $\Phi_G$  is “squashed between” the (equal) minimax and maximin values of the  $\Phi_A$  game. □

## 5 Investing on the Stock Market

In this section we illustrate the extended  $\mu$ -calculus with a small case study [6] which we have modified here to illustrate the new type of fixed point.

An Investor  $I$  has been given the right to make an investment in “futures,” a fixed number of shares in a specific company that he can *reserve* on the first day of any month he chooses. Exactly one month later, the shares will be *delivered* and will collectively have a market *value* on that day.

His problem is to decide when to make his reservation so that his overall profit is maximised, where the profit is the difference between the shares’ price he pays when he makes the reservation and their actual value when he receives them one month later.

The details are as follows:

- (a) The market value (share price)  $v$  of the shares is a whole number of dollars between \$0 and \$10 inclusive; it has a probability  $p$  of going up by \$1 in any month, and  $1-p$  of going down by \$1 — but it remains within those bounds. The probability  $p$  represents short-term market uncertainty.
- (b) Probability  $p$  itself varies month-by-month in steps of 0.1 between zero and one: when  $v$  is less than \$5 the probability that  $p$  will rise is  $2/3$ ; when  $v$  is more than \$5 the probability of  $p$ ’s falling is  $2/3$ ; and when  $v$  is \$5 exactly the probability is  $1/2$  of going either way. The movement of  $p$  represents the Investors’ knowledge of long-term “cyclic second-order” trends.
- (c) There is a cap  $c$  on the value of  $v$ , initially \$10, which has probability  $1/2$  of falling by \$1 in any month; otherwise it remains where it is. The “falling cap” models the fact that the company is in a slow decline.
- (d) If in a given month the Investor does not reserve, then at the very next month he might find he is temporarily *barred* from doing so. But he cannot be barred two months consecutively.
- (e) If he *never* reserves, then he receives no shares and his profit is thus zero.

If it were not for Item (c), the Investor’s strategy would be obvious, that is “wait until  $p = 1$  — however long that takes — and make a reservation

only then.” That way the shares would be certain to increase in the month following his investment and, whatever the (current) market price, the actual price would be greater, guaranteeing him a maximum profit. With (c) however the shares’ value might be forced down to zero before he makes his reservation; then he would make no profit at all.

To analyse this example, we formalise it as a game played by the Market and the Investor: whereas the Investor wishes to maximise his profit, the Market works to minimise it. The utility of the *SSG* is that we can use it to construct a formula describing the fluctuation of the Market as well as the options the Investor has for “playing” it.

First we describe the movement of the Market during one month’s activity (refer to Items (a), (b) and (c) above). The state space is  $(v, p, c)$ , whose value varies according to the transition system

$$\begin{aligned}
 m \quad \hat{=} \quad & v := (v + 1) \sqcap c \quad p \oplus (v - 1) \sqcup 0; \\
 & \mathbf{if} \quad v < 5 \quad \mathbf{then} \quad p := (p - 0.1) \sqcup 0 \quad \textstyle\frac{1}{3} \oplus (p + 0.1) \sqcap 1 \\
 & \mathbf{elseif} \quad v > 5 \quad \mathbf{then} \quad p := (p - 0.1) \sqcup 0 \quad \textstyle\frac{2}{3} \oplus (p + 0.1) \sqcap 1 \\
 & \mathbf{else} \quad \quad \quad p := (p - 0.1) \sqcup 0 \quad \textstyle\frac{1}{2} \oplus (p + 0.1) \sqcap 1 \\
 & \mathbf{fi}; \\
 & c := (c - 1) \sqcup 0 \quad \textstyle\frac{1}{2} \oplus c \quad .
 \end{aligned}$$

The operators  $\sqcup$  and  $\sqcap$  are used conveniently to encode the bounds on the share price and the probability.

Next we combine the (above) transition system with the actions of the Investor and the barring process to describe the complete Investor/Market system. The result is the formula (5) below, the details of which we now explain. We use the label “month” to denote the transition system  $m$  given above, and a constant expectation “Sold” to denote the function returning just the value of the  $v$  component of the state. We then express the *expected* value of the shares one month later as  $\{\text{month}\}\text{Sold}$ , which when interpreted over  $m$  averages Sold over the probabilities of the new share price. For example if  $v$  is 3, and  $p$  is  $3/4$  and  $c$  is 5 at the beginning of the month, then the expected value of the shares one month later is  $3/4 \times 4 + 1/4 \times 2 = 3.5$ ; this is precisely the meaning of  $\{\text{month}\}\text{Sold}$  evaluated at state  $(3, 3/4, 5)$ .

The function “Profit” is defined to be  $\{\text{month}\}\text{Sold} - \text{Sold}$ , and represents the Investor’s expected profit if he decides to reserve in some particular month. If he decides against investing, then his only other option is to wait, and what happens then is expressed by the subformula  $\{\text{month}\}(X \sqcap \{\text{month}\}X)$ ; notice how this introduces the possibility of barring by means of the  $\sqcap$  choice under the control of the Market. Collecting everything together we obtain the

formula describing precisely the Investor/Market game:

$$(5) \quad \textit{Game} \hat{=} (\tau_0 X \cdot \textit{Profit} \sqcup \{\textit{month}\}(X \sqcap \{\textit{month}\}X)) ,$$

where in the subformula  $\{\textit{month}\}(X \sqcap \{\textit{month}\}X)$ , the variable  $X$  has been bound to a fixed point  $\tau_0$ . We have chosen  $\tau_0$  since we are equating infinite paths with the Investor’s decision never to invest, and his profit in that situation is, of course 0 (refer to Item(e)). Note, however, that this is not a standard *SSG*, nor is  $\tau_0$  simply a least- or greatest fixed point in disguise. To see that, observe that **Profit** can take on both positive and negative values, depending on whether the share price is more likely to increase or to decrease at the beginning of a particular month. Thus 0 is neither the greatest nor the least value within the complete range of payoffs.

To compute the value of Formula (5) (and therefore the Investor’s expected overall payoff) our theorems suggest first to translate it into a game with payoffs in the range  $[0, 1]$  using an affine transformation (Lem. 3.1); next to convert that game into a standard *SSG* (Thm. 4.8), next to solve that game (Thm. 2.1) and then to translate the result back again (Lem. 3.1). Unfortunately this is quite inefficient due to the introduced double fixed point (especially if iterative methods must be used).

However in this case there is an efficient solution. We notice that in this case the value can be computed by iterating from  $\underline{0}$ .

**Lemma 5.1** *In the game  $(\tau_x X \cdot \Phi)$ , if  $\underline{x} \leq F(\underline{x})$  then  $V^* = \lim_{n \rightarrow \infty} F^n(\underline{x})$ , where expectation-to-expectation function  $F$  is the denotation of  $\Phi(\cdot)$  as a function of the value denoted by  $X$ .*

**Proof.** Under the assumption given, the value of the  $x$ -gadget game (Def. 4.3) reduces easily to  $\lim_{n \rightarrow \infty} F^n(\underline{x})$  in the case that the payoffs are all in  $\mathcal{ES}$ ; the general case follows from Lem. 3.1. □

Using Lem. 5.1 we created a *Mathematica* script which encoded the game given by (5), and iterated directly from 0. The initial value of  $p$  was set at 0.5 and the cap  $c$  to the maximum 10; a selection of the results appear in the table below:

Initial share value	0	1	2	3	4	5	6	7	8	9	10
Best expected profit	0.59	0.57	0.53	0.47	0.40	0.33	0.27	0.23	0.20	0.19	0.18

The results show that there is always some probability of observing a rising market, but that waiting longer for a more favourable value of  $p$  is possible when the initial share price is far from the value of the cap  $c$ .

## 6 Conclusions and further work

In this paper we have introduced a novel quantitative two-player game which models situations where neither player wins decisively. We defined a new kind of fixed point as a convenient way to model this situation directly, but showed how the original  $qM\mu$  can in fact model it using interleaved fixed points. The consequence of encoding it as a standard *SSG* suggests the possibility of developing algorithms for computing the value using some of the techniques developed for stochastic parity games [1]; however we note that the increase in alternation depth implicit in the gadget-encoding of Def. 4.3 might impact their efficiency. That remains a topic for further investigation.

## References

- [1] Chatterjee, K., M. Jurdzinski and T. A. Henzinger, *Quantitative stochastic parity games*, in: *Proceedings of the 15th Annual Symposium on Discrete Algorithms*, ACM/SIAM, 2004, pp. 114–123.
- [2] de Alfaro, L. and R. Majumdar, *Quantitative solution of omega-regular games*, *Journal of Computer and System Sciences* **68** (2004), pp. 374–397.
- [3] Grimmett, G. and D. Welsh, “Probability: an Introduction,” Oxford Science Publications, 1986.
- [4] Huth, M. and M. Kwiatkowska, *Quantitative analysis and model checking*, in: *Proceedings of 12th annual IEEE Symposium on Logic in Computer Science*, 1997.
- [5] Kozen, D., *Results on the propositional  $\mu$ -calculus*, *Theoretical Computer Science* **27** (1983), pp. 333–354.
- [6] McIver, A. and C. Morgan, *Games, probability and the quantitative  $\mu$ -calculus  $qM\mu$* , in: *Proc. LPAR*, LNAI **2514** (2002), pp. 292–310, revised and expanded at [8]; also appears at [7, Chap. 11].
- [7] McIver, A. and C. Morgan, “Abstraction, Refinement and Proof for Probabilistic Systems,” *Technical Monographs in Computer Science*, Springer, New York, 2005.
- [8] McIver, A. and C. Morgan, *Results on the quantitative  $\mu$ -calculus  $qM\mu$*  (2005), to appear in *ACM TOCL*.
- [9] Morgan, C. and A. McIver, *A probabilistic temporal calculus based on expectations*, in: L. Groves and S. Reeves, editors, *Proc. Formal Methods Pacific '97* (1997), available at [11]; also appears at [7, Chap. 9].
- [10] Morgan, C. and A. McIver, *An expectation-based model for probabilistic temporal logic*, *Logic Journal of the IGPL* **7** (1999), pp. 779–804, also appears at [7, Chap.10].
- [11] PSG, *Probabilistic Systems Group: Collected reports*, <http://web.comlab.ox.ac.uk/oucl/research/areas/probs/bibliography.html>.
- [12] Stirling, C., *Local model checking games*, in: *CONCUR 95*, number 962 in LNCS (1995), pp. 1–11, extended abstract.
- [13] Stoy, J., “Denotational Semantics: The Scott-Strachey Approach to Programming Language Theory,” MIT Press, 1977.
- [14] Vardi, M., *Automatic verification of probabilistic concurrent finite-state programs*, in: *Proc. 26th IEEE Symp. on Foundations of Computer Science*, Portland, 1985, pp. 327–338.