A Testing Equivalence for Reactive Probabilistic Processes

M.Z. Kwiatkowska and G.J. Norman

School of Computer Science,
University of Birmingham,
Edgbaston, Birmingham B15 2TT, UK

Abstract

We consider a generalisation of Larsen and Skou's [19] reactive probabilistic transition systems which exhibit three kinds of choice: action-guarded probabilistic choice, external (deterministic) and internal (non-deterministic) choice. We propose an operational preorder and equivalence for processes based on testing. Milner's button pushing experiments scenario is extended with random experiments by assessing the *probability* of processes passing a test. Two processes are then identified with respect to the testing equivalence if they pass all tests with the same probability. The derived equivalence is a congruence for a subcalculus of CSP extended with action-guarded probabilistic choice. It is coarser than probabilistic bisimulation, yet non-probabilistic branching-time, and differs from probabilistic equivalences developed for CSP [20,22,26]. We provide a logical characterization of the equivalence in terms of the quantitative interpretation of HML of [14] and show how fixed points can be added to the logic.

1 Introduction

Many probabilistic extensions of process algebras have been proposed to date, such as those based on CCS [21], CSP [7] and ACP [4]. Probabilistic bisimulation, introduced by Larsen and Skou [19] for reactive systems and extended with non-determinism and time by Hansson [12], is a generalisation of Milner's bisimulation. Other probabilistic process equivalences include probabilistic simulation of Segala and Lynch [25], Wang Yi and Larsen's testing equivalence [27], and CSP equivalences of Morgan et al. [22], Lowe [20] and Seidel [26].

Probabilistic bisimulation is strongly related to bisimulation [6] and has many useful properties: it has a logical characterization in terms of the Hennessy-Milner logic [19] (see also [8]), has an efficient (polynomial) decision procedure [2], and is a congruence for typical process operators. For example, van Glabbeek et al. [11] show that probabilistic bisimulation is a congruence

This is a preliminary version. The final version can be accessed at
URL: http://www.elsevier.nl/locate/entcs/volume16.html

over their calculus PCCS (which contains all the usual SCCS operators) and Baier and Kwiatkowska [3] show congruence properties of full CCS extended with action-guarded probabilistic choice. Generally, if one works with a fine (or strong) equivalence such as probabilistic bisimulation then almost all CCS or CSP operators can be adapted to the probabilistic setting. However, there are cases when probabilistic bisimulation is too fine, as it discriminates between processes that cannot be distinguished under a realistic testing scenario.

One alternative is to work with a weaker (or coarser) equivalence, i.e. one that only distinguishes processes that can be differentiated by external observations. The difficulty with this approach is that only a subset of operators can be considered if we wish to ensure our equivalence is a congruence (see e.g. Jou and Smolka [17], where even restriction forces both trace and failure equivalence to fail to be congruences, and also in [26] and [20] where hiding cannot be defined); the latter is an important property, since without it any resulting denotational model will not be compositional.

This paper is motivated by the need to work with a process language which allows both external choice (determined by the environment) and internal choice (determined by the process itself) in the sense of CSP [24], together with probabilistic choice (which we shall assume to be action-guarded). In addition, we shall assume that the probabilistic choice is *internal*[1] , that is, the point at which probabilistic choices occur is unimportant, since they are made neither by the process nor by the environment, but by some prescribed probability distribution. An example of such a situation involving *scratch cards* can be found in [22]. In such cases probabilistic bisimulation is too fine, as it makes distinctions between processes that cannot be distinguished by a reasonable notion of observation. To illustrate this point, consider the processes given in Figure 1 below.

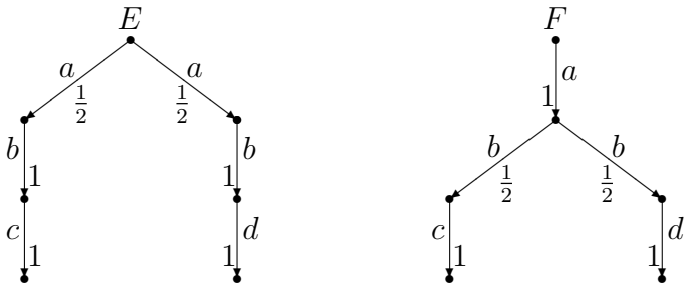

Fig. 1. Probabilistic bisimulation is too fine.

Both E and F make a random choice with probability $\frac{1}{2}$ at some stage in their computations. Each random choice can be thought of as flipping a coin and then selecting the left or the right branch depending on the outcome. The process E, therefore, first flips a coin and then performs the actions abc if the coin lands on heads, or the actions abd if the coin lands on tails. The process

[1] There may be cases when *external* probabilistic choice is needed, which we shall not discuss in this paper.

F, on the other hand, first performs an a action, and then flips a coin and performs the actions bc if the coin lands on heads, or the actions bd if the coin lands on tails. Since *performing an a action before or after flipping a coin has no effect on whether the coin lands on heads or tails*, the processes should be observationally equivalent. However, probabilistic bisimulation will distinguish between these processes, due to the difference in their *probabilistic branching* behaviour.

In this paper we consider an extension of Larsen and Skou's reactive systems [19] so that they allow three types of choice: probabilistic, external and internal. For such systems we formulate an equivalence based on testing, developed in the setting of Milner's button pushing experiments by the addition of *random experiments*. These are experiments in which *the probability of a given process passing a test or not* can be assessed. One process is then said to be greater than or equal to another if it passes all the tests with probability at least as high as the former. Two processes are equivalent if they pass all tests with the same probability. The idea for this equivalence is to only make distinctions that are in some sense observable (for which there is a test that the processes pass with different probabilities), while at the same time ensuring that it is a congruence for the three types of choice operators, particularly external choice. We compare the derived equivalence with similar probabilistic equivalences.

Finally, we provide a 'logical' characterization in terms of the quantitative interpretation of the Hennessy-Milner logic introduced by Huth and Kwiatkowska [14]. Each formula of the quantitative `HML` is interpreted as a map from processes to the interval [0,1], giving the probability of the process satisfying this formula. Each such formula is shown to correspond to a random experiment, and vice-versa. We also show how a fixed point operator can be added to the quantitative `HML`.

For full details of the results presented here the reader is referred to [23]. A fully abstract metric-space denotational semantics for the calculus presented here can be found in [23,18].

2 A Testing Scenario with Random Experiments

In this section, we first introduce *reactive probabilistic transition systems* which extend Larsen and Skou's probabilistic labelled transition systems [19] by allowing processes of the system to exhibit three types of choice: (internal action-guarded) probabilistic, external (deterministic) and internal (nondeterministic). Next we introduce an operational preorder over reactive probabilistic transition systems based on testing which will distinguish two processes only if they have observably different behaviour. We then compare this order and equivalence with alternative probabilistic equivalences and give suitable examples.

2.1 Reactive Probabilistic Processes

Let D be a set. A (discrete) *probability distribution* on D is a function $\pi : D \longrightarrow [0,1]$ such that $\sum_{d \in D} \pi(d) = 1$. Furthermore, let $\mu(D)$ denote the set of discrete probability distributions on D. Let A and S be sets. A subset X of $A \times S$ is said to satisfy the *reactiveness condition* if, for any distinct $(a_1, s_1), (a_2, s_2) \in X$: $a_1 \neq a_2$. Furthermore, let $\mathcal{P}_{\rm fr}(\cdot \times \cdot)$ denote the powerset operator restricted to only finite reactive subsets of cartesian products satisfying the reactiveness condition.

Definition 2.1 A *Reactive Probabilistic Transition System* is a tuple $(\mathcal{R}, \mathcal{A}ct, \rightarrow)$, where $\mathcal{R}$ is a set of states, $\mathcal{A}ct$ is a finite set of actions and $\rightarrow$ a transition relation

$$\rightarrow \subseteq \mathcal{R} \times \mathcal{P}_{\rm fr}(\mathcal{A}ct \times \mu(\mathcal{R}))$$

satisfying: for all $E \in \mathcal{R}$ there exists $S \in \mathcal{P}_{\rm fr}(\mathcal{A}ct \times \mu(\mathcal{R}))$ such that $(E, S) \in \rightarrow$. We write $E \rightarrow S$ instead of $(E, S) \in \rightarrow$.

Note that elements $E \in \mathcal{R}$ of a reactive probabilistic transition system (the processes) are associated via the (unlabelled) transition relation $\rightarrow$ with reactive *sets* of pairs consisting of an action $a \in \mathcal{A}ct$ and a probability distribution π on the processes $\mathcal{R}$. Intuitively, any such $S = \{(a_1, \pi_1), \ldots, (a_m, \pi_m)\}$ should be viewed as a reactive probabilistic process *deterministic* on its first step, which offers to the environment the menu of actions $a_1, \ldots, a_m$, and after a_i for some $1 \leq i \leq m$ has been selected, the process continues according to the distribution π_i, that is, the probability of behaving as F is given by $\pi_i(F)$. Uniqueness of this distribution is guaranteed by the reactiveness condition. The case when $S = \emptyset$, the inactive process, is allowed. These 'deterministic' probabilistic processes are, in fact, equivalent to Larsen and Skou's probabilistic transition systems [19]; using their terminology, for any $S \in \mathcal{P}_{\rm fr}(\mathcal{A}ct \times \mu(\mathcal{R}))$:

$$S \stackrel{a}{\longrightarrow}_{\lambda} F \text{ if and only if } (a, \pi) \in S \text{ for some } \pi \in \mu(\mathcal{R}) \text{ and } \pi(F) = \lambda.$$

Non-determinism is introduced by allowing a choice between deterministic processes: for any $E \in \mathcal{R}$ and distinct $S_1, S_2 \in \mathcal{P}_{\rm fr}(\mathcal{A}ct \times \mu(\mathcal{R}))$, if $E \rightarrow S_1$ and $E \rightarrow S_2$, then E makes a *non-deterministic choice* between continuing as the process S_1 or S_2. The class of all reactive systems thus exhibits (internal action-guarded) probabilistic, deterministic and non-deterministic choice.

2.2 Testing Reactive Probabilistic Processes

We develop an operational preorder on reactive probabilistic transition systems with the help of the testing language `T` which we now introduce. The testing scenario is based on Milner's button-pushing experiments on transition systems [21], where we suppose we have a series of buttons, one for every action ($a \in \mathcal{A}ct$), which an experimenter can press one at a time. The process will remain in rest if no buttons are pressed (we denote this experiment

by ω), and otherwise it will react by performing the corresponding action, in which case the button will go down and the experiment *succeeds*, or refusing to perform the action, in which case the button will not go down and the experiment *fails*. This scenario is extended to handle random experiments as follows. First we define two experiments t_i and t_j to be *independent* if the first steps of the experiments are associated with pressing different buttons.

Next we introduce three experiments to respectively capture the three types of behaviour exhibited by processes of a reactive probabilistic transition system (probabilistic, non-deterministic and external).

(i) $a.t$, where $a \in \mathcal{A}ct$: this experiment corresponds to pushing the a-button and then, if the button goes down, performing the experiment t.

(ii) $(\!|t|\!)$: this experiment corresponds to making sufficiently many copies of the process being tested, so that any non-deterministic choice the process can make will occur on at least one of the copies made, and then performing the experiment t on *each* of the copies.

(iii) $((\!|t_1|\!), \ldots, (\!|t_m|\!))$ and $[a_1.t_1, \ldots, a_m.t_m]$, where for all $1 \leq i \neq j \leq m$ the experiments $(\!|t_i|\!)$ and $(\!|t_j|\!)$, and the experiments $a_i.t_i$ and $a_j.t_j$, are independent: the experiment $((\!|t_1|\!), \ldots, (\!|t_m|\!))$ corresponds to making m copies of the process being tested and then performing the experiment $(\!|t_i|\!)$ on *one* of the copies for all $1 \leq i \leq m$ ($[a_1.t_1, \ldots, a_m.t_m]$ is similar).

Intuitively, the success or failure of a process passing an experiment corresponds to the success or failure of *one run* (or execution) of the process being experimented on, under different conditions. This motivates the construction $(\!|t|\!)$ and also the constructions $((\!|t_1|\!), \ldots, (\!|t_m|\!))$ and $[a_1.t_1, \ldots, a_m.t_m]$ (and the restrictions of independence we have imposed on them). First, $(\!|t|\!)$ corresponds to the changes the *demons* introduce to influence the non-deterministic choices that processes make. We note that through the construct $(\!|t|\!)$ we require that the demons must make all non-deterministic choices that a process can make possible within a finite period. Second, $((\!|t_1|\!), \ldots, (\!|t_m|\!))$ (similarly for $[a_1.t_1, \ldots, a_m.t_m]$) corresponds to changes in the environment, that is, changes in the actions the processes are allowed to perform, which we accomplish by making copies of processes and then pressing *different* buttons on each of these copies. To ease notation, when forming the test $((\!|t_1|\!), \ldots, (\!|t_m|\!))$ (similarly for $[a_1.t_1, \ldots, a_m.t_m]$) we require that $(\!|t_i|\!)$ and $(\!|t_j|\!)$ are independent for all $1 \leq i \neq j \leq m$. We note that copying was first introduced by Abramsky [1] for non-probabilistic processes, and by Larsen and Skou [19] for probabilistic processes.

If we now perform these tests on processes of a reactive probabilistic transition system $(\mathcal{R}, \mathcal{A}ct, \rightarrow)$, we will be performing *random experiments*, since for any $E \in \mathcal{R}$ the success of the experiment will depend on the probabilistic choices within the process. Moreover, the test $(\!|r|\!)$ will, in fact, give rise to a *set* of probabilities, each one corresponding to the probability of the E passing the test r when one of its possible non-deterministic choices is made.

As a result, we will be unable to calculate the *exact* probability of processes passing tests and instead we will only calculate the *greatest lower bound* and the *least upper bound* on the probability of E passing r. These are the only two realistic options, as there is no way of calculating any meaningful average, since the choices are non-deterministic and so we are unable to estimate the frequency of each such choice being made.

The syntax of our testing language $\mathtt{T}_\omega$ is as follows.

Definition 2.2 Let $\mathtt{T}$ and $\mathtt{T}_\omega$, with elements t and T respectively, be the testing languages defined inductively as follows:

$$r ::= \omega \mid [a.T, \ldots, a.T]$$
$$t ::= (\!|r|\!)$$
$$T ::= (t, \ldots, t)$$

where $a \in \mathcal{A}ct$.

To capture the outcome of random experiments as described above we define the maps $\mathsf{R}_{\mathbf{glb}}$ and $\mathsf{R}_{\mathbf{lub}}$ from $\mathcal{R}$ and $\mathtt{T}_\omega$ to the unit interval which, for any process $E \in \mathcal{R}$ and test $(\!|r|\!) \in \mathtt{T}$, yield the greatest lower bound and the least upper bound on *the probability of E passing the test r* respectively. We mention that intervals were also used in [27].

Definition 2.3 Let $\mathsf{R}_{\mathbf{glb}}, \mathsf{R}_{\mathbf{lub}} : \mathcal{R} \longrightarrow (\mathtt{T}_\omega \longrightarrow [0,1])$ be the maps defined inductively on $\mathtt{T}_\omega$ where R_* stands for either $\mathsf{R}_{\mathbf{glb}}$ or $\mathsf{R}_{\mathbf{lub}}$. For any $E \in \mathcal{R}$ put:

$$\mathsf{R}_{\mathbf{glb}}(E)((\!|r|\!)) = \min_{E \to S} \mathsf{R}_{\mathbf{glb}}(S)(r), \quad \mathsf{R}_{\mathbf{lub}}(E)((\!|r|\!)) = \max_{E \to S} \mathsf{R}_{\mathbf{lub}}(S)(r)$$
$$\text{and} \quad \mathsf{R}_*(E)((t_1, \ldots, t_m)) = \prod_{j=1}^{m} \mathsf{R}_*(E)(t_j)$$

where for any $S \in \mathcal{P}_{\mathrm{fr}}(\mathcal{A}ct \times \mu(\mathcal{R}))$ and $1 \leq i \leq m$ put:

$$\mathsf{R}_*(S)(\omega) = 1, \quad \mathsf{R}_*(S)([a_1.T_1, \ldots, a_m.T_m]) = \prod_{i=1}^{m} \mathsf{R}_*(S)(a_i.T_i) \quad \text{and}$$

$$\mathsf{R}_*(S)(a.T) = \begin{cases} \sum_{F \in \mathcal{R}} \pi(F) \cdot \mathsf{R}_*(F)(T) & \text{if } (a, \pi) \in S \text{ for some } \pi \in \mu(\mathcal{R}) \\ 0 & \text{otherwise.} \end{cases}$$

The intuition for how the values of $\mathsf{R}_{\mathbf{glb}}(E)((\!|r|\!))$ and $\mathsf{R}_{\mathbf{lub}}((\!|r|\!))$ are calculated should be clear from the discussion above and the finiteness assumptions. Next, $\mathsf{R}_*(S)(\omega)$ calculates the probability of S passing the test ω, and since any process can pass ω (as no buttons are pressed) we set this value to 1. In the case for $\mathsf{R}_*(E)((t_1, \ldots, t_m))$ we want to calculate the probability of E or S passing *all* the tests $t_1, \ldots, t_m$, which we achieve by multiplying the probabilities of E or S passing each test t_i. Multiplication can be used since by construction, for any $1 \leq j \neq k \leq m$ and for any $1 \leq i \neq j \leq m$, the tests t_j and t_k correspond to pressing different buttons at their first stage, hence corresponding

to different probability distributions, and therefore the probabilities of these tests being passed are *independent*. The case of $\mathsf{R}_*(S)([a_1.T_1, \ldots, a_m.T_m])$ is similar. Finally, $\mathsf{R}_*(S)(a.T)$ calculates the probability of the 'deterministic' probabilistic process S performing paths which have the initial action a and then pass the test T: we set $\mathsf{R}_*(S)(a.T) = 0$ if S cannot perform the action a, that is $(a, \pi) \notin S$ for any $\pi \in \mu(\mathcal{R})$, and to the *weighted sum* of $\pi(F) \cdot \mathsf{R}_*(F)(T)$ over all $F \in \mathcal{R}$ otherwise.

Since by definition $\mathsf{R}_{\mathbf{glb}}(E)((\!|\omega|\!)) = \mathsf{R}_{\mathbf{lub}}(E)((\!|\omega|\!)) = 1$ for all reactive probabilistic transition systems $(\mathcal{R}, \mathcal{A}ct, \rightarrow)$ and $E \in \mathcal{R}$, to simplify the notation we will denote any occurrence of the test $(\!|\omega|\!)$ by ω.

With the help of the above maps, we are now in a position to define our operational order and subsequent equivalence on all reactive probabilistic transition systems. We simply require that the process higher up the order must pass all tests with probability *at least as high* as those below.

Definition 2.4 For any $E, F \in \mathcal{R}$, $E \sqsubseteq^{\mathbf{glb}} F$ if $\mathsf{R}_{\mathbf{glb}}(E)(T) \leq \mathsf{R}_{\mathbf{glb}}(F)(T)$ and $E \sqsubseteq^{\mathbf{lub}} F$ if $\mathsf{R}_{\mathbf{lub}}(E)(T) \leq \mathsf{R}_{\mathbf{lub}}(F)(T)$ for all $T \in \mathtt{T}_\omega$ respectively. Moreover, for any $E, F \in \mathcal{R}$, $E \sqsubseteq^{\mathsf{R}} F$ if $E \sqsubseteq^{\mathbf{glb}} F$ and $E \sqsubseteq^{\mathbf{lub}} F$, and $E \overset{\mathsf{R}}{\sim} F$ if $E \sqsubseteq^{\mathsf{R}} F$ and $F \sqsubseteq^{\mathsf{R}} E$.

The following lemma illustrates why we need only consider the tests $\mathtt{T}$ as opposed to the (larger) set of tests $\mathtt{T}_\omega$.

Lemma 2.5 *For any $E, F \in \mathcal{R}$:*

(i) $E \sqsubseteq^{\mathbf{glb}} F$ *if and only if* $\mathsf{R}_{\mathbf{glb}}(E)(t) \leq \mathsf{R}_{\mathbf{glb}}(F)(t)$ *for all* $t \in \mathtt{T}$.

(ii) $E \sqsubseteq^{\mathbf{lub}} F$ *if and only if* $\mathsf{R}_{\mathbf{lub}}(E)(t) \leq \mathsf{R}_{\mathbf{lub}}(F)(t)$ *for all* $t \in \mathtt{T}$.

We note that our testing scenario differs from that of Larsen and Skou's in that their testing scenario removes the syntactic restriction of independence we impose on the construct $(t, \ldots, t)$. As a result, the two approaches attach a different meaning to the phrase "the probability of a process passing a test". In our approach, the probability of a process passing a test corresponds to the probability of *one run* (or execution) of the process passing a test, with the addition that we allow the value to correspond to the probability of the same run of a process passing a test under different conditions, for example due to changes in the behaviour of the environment. Thus, in our setting, the probability of a process passing the test $(a.t, b.t)$ is the probability of some run of the process passing the test $a.t$ when the environment offers the action a, and the *same* run passing the test $b.t$ when the action b is offered. It should therefore come as no surprise that our testing equivalence does not coincide with the equivalence introduced in [19].

2.3 A Comparison with Alternative Equivalences on Probabilistic Processes

We now compare our operational order $\sqsubseteq^{\mathsf{R}}$ with probabilistic equivalences known from the literature. One such equivalence is Larsen and Skou's prob-

abilistic bisimulation [19], which turns out to be finer than our equivalence. To see this, consider again the processes in Figure 1, which are distinguished by probabilistic bisimulation. Analysing the values of maps $\mathsf{R}_{\mathbf{glb}}(E)$, $\mathsf{R}_{\mathbf{lub}}(E)$, $\mathsf{R}_{\mathbf{glb}}(F)$ and $\mathsf{R}_{\mathbf{lub}}(F)$ over all tests $\mathtt{T}$, we see that the maps agree. The result is summarised in the table below, with the zero values omitted:

$[a.\omega]$	$[a.(([b.\omega]))]$	$[a.(([b.(([c.\omega]))]))]$	$[a.(([b.(([d.\omega]))]))]$
1	1	1/2	1/2

Thus, the order $\sqsubseteq^{\mathsf{R}}$ will not distinguish between the processes E and F, and hence $E \overset{\mathsf{R}}{\sim} F$.

Other equivalences that are finer than $\sqsubseteq^{\mathsf{R}}$ and also distinguish the processes given in Figure 1, and which we therefore view as too fine, include: Hansson's extension of probabilistic bisimulation to a model allowing nondeterminism [12], Segala and Lynch's probabilistic simulation [25] and Wang Yi and Larsen's testing equivalence [27]. Nevertheless, when considering models with *external* probabilistic choice, the probabilistic branching structure may become important, since the probabilistic choices the processes make will depend on the choices made by the environment.

Of equivalences coarser than $\overset{\mathsf{R}}{\sim}$, which would therefore identify any two processes that $\overset{\mathsf{R}}{\sim}$ finds equivalent, and in particular could not distinguish between the processes of Figure 1, there are several based on extending traces, failures and readies [7]. Such extensions are based on incorporating the probabilities of processes performing a trace, and then refusing or accepting to then perform a certain set of actions, and include equivalences formulated by Seidel [26], Lowe [20] and Jou and Smolka [17]. In our reactive setting, these equivalences are too coarse: although they do capture the probabilistic behaviour of processes, they are linear-time based equivalences [10], and therefore do not capture the branching behaviour associated with choices other than probabilistic, such as deterministic choice which reactive systems admit. To illustrate this consider the processes in Figure 2 below which are distinguished by our equivalence $\overset{\mathsf{R}}{\sim}$ (and hence probabilistic bisimulation) but are equivalent under probabilistic extensions of trace, failure and ready equivalences.

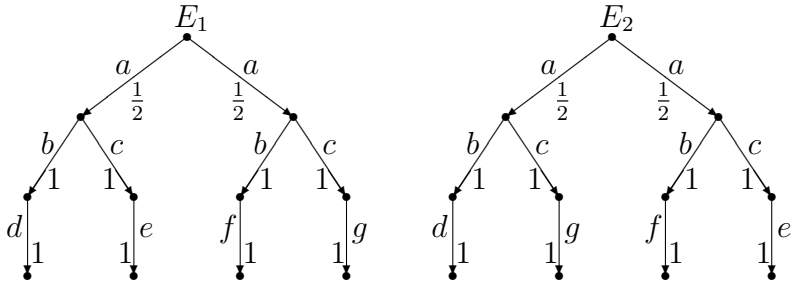

Fig. 2. Trace, failure and ready equivalence are too coarse.

Observe that E_1 can reach an intermediate state (after performing the action a with probability $\frac{1}{2}$) where there is an *external choice* between performing a

b transition followed by a d transition, and performing a c transition followed by an e transition. In contrast, E_2 cannot reach such a state.

The final equivalence we mention, neither finer nor coarser than $\overset{\mathsf{R}}{\sim}$, is that introduced by Morgan et al. [22]. Although, similarly to [20], the equivalence of [22] is based on the failures model of CSP, it is essentially different in that it is instead based on how processes "make decisions", and more precisely on what the process "is". The latter is achieved by intuitively considering "the probability that probabilistic processes are standard CSP processes". We feel that their equivalence is too fine in certain cases. As an example consider the processes E_3 and E_4 given in Figure 3 below.

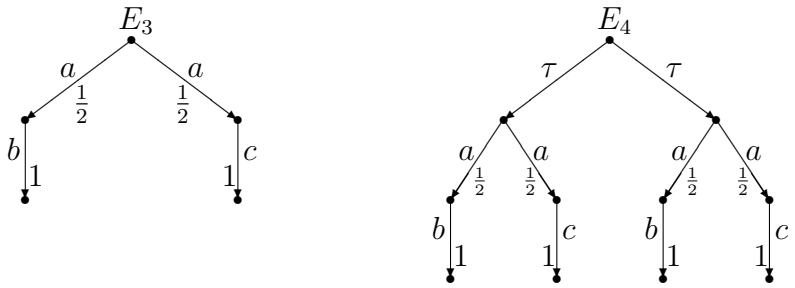

Fig. 3. Morgan et al.'s equivalence is too fine.

First, observe that E_3 can either perform the trace ab or the trace ac, both with probability $\frac{1}{2}$. Moreover, no matter which internal choice E_4 can make, the outcome will match the behaviour of E_3. Therefore, these processes should be observationally equivalent. However, in the approach of Morgan et al. [22], the processes are distinguished: for example, the probability that E_3 is the CSP process $a \rightarrow (b \rightarrow \mathbf{0})$ is $\frac{1}{2}$, whereas the probability that E_4 is the process $a \rightarrow (b \rightarrow \mathbf{0})$ is $\frac{1}{4}$ since E_4 only becomes the process $a \rightarrow (b \rightarrow \mathbf{0})$ when both instances of E_3 in E_4 choose to perform the trace ab.

We end this section by giving examples of reactive probabilistic processes which illustrates that the order $\sqsubseteq^{\mathsf{R}}$ is non-probabilistic branching time.

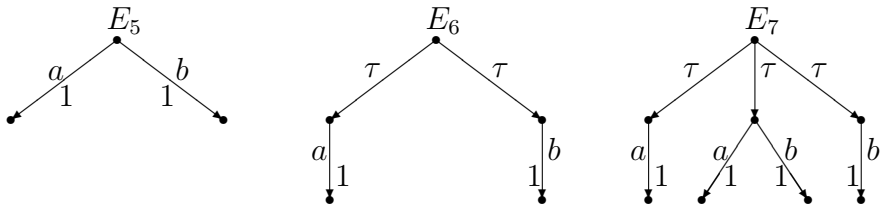

Fig. 4. $\sqsubseteq^{\mathsf{R}}$ is non-probabilistic branching time ($E_4 \sqsubseteq^{\mathsf{R}} (\overset{\mathsf{R}}{\not\sim}) E_5 \sqsubseteq^{\mathsf{R}} (\overset{\mathsf{R}}{\not\sim}) E_3$).

3 Logical Semantics

In this section we give a logical semantics to reactive probabilistic transition systems using Hennessy-Milner Logic (HML) [13]. We adapt and extend Huth and Kwiatkowska's non-standard interpretation [14] for HML over processes of

Larsen and Skou's probabilistic transition systems [19] to our reactive probabilistic transition systems.

We begin by recalling the logic HML (with finite conjunctions) together with the non-standard interpretation of Huth and Kwiatkowska [14] originally introduced for Larsen and Skou's probabilistic transition systems. We omit additional binary operators considered in [14], for example disjunction ($\vee$), and fixed point operators.

Definition 3.1 (Hennessy-Milner Logic [13]) The logic HML is defined inductively on the syntax:

$$\phi ::= \mathtt{true} \mid \langle a \rangle \phi \mid \neg \phi \mid \phi \wedge \phi$$

where a ranges over a set of actions $\mathcal{A}ct$.

Definition 3.2 (c.f. [14]) Let $[\![.]\!] : \mathtt{HML} \longrightarrow (P \longrightarrow [0,1])$ be the map defined inductively on formulae of HML for any process E of Larsen and Skou's probabilistic transition systems [19] as follows:

$$\begin{aligned} [\![\mathtt{true}]\!]E &= 1 \\ [\![\langle a \rangle \phi]\!]E &= \textstyle\sum_{E \xrightarrow{a}_{\lambda} F} \lambda \cdot [\![\psi]\!]F \\ [\![\neg \phi]\!]E &= 1 - [\![\phi]\!]E \\ [\![\phi_1 \wedge \phi_2]\!]E &= [\![\phi_1]\!]E \cdot [\![\phi_2]\!]E. \end{aligned}$$

Observe that conjunctions are interpreted as multiplication, and thus we must impose independence of the corresponding events to ensure probabilistic soundness. We achieve this through a syntactic restriction of disjointness of the initial actions of the conjuncts (see the mapping act below). This can be extended to the full HML with the help of conditional probabilities.

Definition 3.3 Let $\mathrm{act} : \mathtt{HML} \longrightarrow \mathcal{P}_{\mathrm{f}}(\mathcal{A}ct)$ be the mapping defined inductively on the syntax of HML as follows:

$$\begin{aligned} \mathrm{act}(\mathtt{true}) &= \emptyset \\ \mathrm{act}(\langle a \rangle \phi) &= \{a\} \\ \mathrm{act}(\neg \phi) &= \mathrm{act}(\phi) \\ \mathrm{act}(\phi_1 \wedge \phi_2) &= \mathrm{act}(\phi_1) \cup \mathrm{act}(\phi_2). \end{aligned}$$

Since the interpretation $[\![\cdot]\!]$ of HML (see Definition 3.2) is given for processes of Larsen and Skou's systems which, as mentioned earlier, are equivalent to the deterministic components of processes of a reactive probabilistic transition system $(\mathcal{R}, \mathcal{A}ct, \rightarrow)$, that is, elements of $\mathcal{P}_{\mathrm{fr}}(\mathcal{A}ct \times \mu(\mathcal{R}))$, we can simply adapt Definition 3.2 to our reactive setting as follows. Let $S \in \mathcal{P}_{\mathrm{fr}}(\mathcal{A}ct \times \mu(\mathcal{R}))$; we

need to replace the clause for $\langle a\rangle\phi$ in Definition 3.2 by:

$$[\![\langle a\rangle\phi]\!]S \stackrel{\text{def}}{=} \begin{cases} \sum_{F\in\mathcal{R}} \pi(F)\cdot[\![\phi]\!]F \text{ if } (a,\pi)\in S \text{ for some } \pi\in\mu(\mathcal{R}) \\ \qquad 0 \qquad\qquad \text{otherwise.} \end{cases}$$

Next, to extend this interpretation of HML to all the processes of $(\mathcal{R},\mathcal{A}ct,\rightarrow)$, we need to incorporate the non-deterministic behaviour. Similarly to the non-probabilistic case, where non-deterministic behaviour is often represented by processes being able to perform hidden (τ) actions (for example, when giving a logical characterisation of weak bisimulation [21]), we add an operator of the form $\langle\varepsilon\rangle\phi$ to the syntax of HML, where for any labelled transition system and process P of the system, $\langle\varepsilon\rangle\phi$ is interpreted as follows:

$$[\![\langle\varepsilon\rangle\phi]\!]P \stackrel{\text{def}}{=} \max\{[\![\phi]\!]Q \mid P\Rightarrow Q\}.$$

Here $P\Rightarrow Q$ holds if there exists a path from P to Q consisting of an arbitrary number (≥ 0) of τ-steps. Intuitively, a process P satisfies the formula $\langle\varepsilon\rangle\phi$, that is, $[\![\langle\varepsilon\rangle\phi]\!]P = 1$, if P *can* make a non-deterministic choice to evolve as a process which will satisfy ϕ. Adapting this to $(\mathcal{R},\mathcal{A}ct,\rightarrow)$, we have the following interpretation of $\langle\varepsilon\rangle\phi$ for any $E\in\mathcal{R}$:

$$[\![\langle\varepsilon\rangle\phi]\!]E \stackrel{\text{def}}{=} \max\{[\![\phi]\!]S \mid E\rightarrow S\}$$

since E makes a non-deterministic choice between behaving as any $S\in\mathcal{P}_{\text{fr}}(\mathcal{A}ct\times\mu(\mathcal{R}))$ such that $E\rightarrow S$ (the reader should note the resemblance to our tests).

We also add the dual of $\langle\varepsilon\rangle\phi$, namely $[\varepsilon]\phi$, where, intuitively, a (non-probabilistic) process P satisfies the formula $[\varepsilon]\phi$ if *all* the processes that P can evolve to by making a non-deterministic choice satisfy ϕ. Since by definition of HML $[.] = \neg\langle.\rangle\neg$, using Definition 3.2 we define the interpretation of $[\varepsilon]\phi$ by: for any $E\in\mathcal{R}$:

$$[\![[\varepsilon]\phi]\!]E \stackrel{\text{def}}{=} \min\{[\![\phi]\!]S \mid E\rightarrow S\}.$$

Furthermore, since any $S\in\mathcal{P}_{\text{fr}}(\mathcal{A}ct\times\mu(\mathcal{R}))$ is a deterministic probabilistic process, we set:

$$[\![\langle\varepsilon\rangle\phi]\!]S \stackrel{\text{def}}{=} [\![\phi]\!]S \text{ and } [\![[\varepsilon]\phi]\!]S \stackrel{\text{def}}{=} [\![\phi]\!]S.$$

We identify the following two subsets of HML, denoted $\text{HML}_{\text{r}}^{\langle\varepsilon\rangle}$ and $\text{HML}_{\text{r}}^{[\varepsilon]}$, where the intuitive meaning of $\text{HML}_{\text{r}}^{\langle\varepsilon\rangle}$ and $\text{HML}_{\text{r}}^{[\varepsilon]}$ is that processes *may* or *must* validate formulae respectively – in the probabilistic sense of course – where the map $\text{act}(\cdot)$ is extended by: $\text{act}(\langle\varepsilon\rangle\phi)\stackrel{\text{def}}{=}\text{act}(\phi)$.

Definition 3.4 The sublanguage $\text{HML}_{\text{r}}^{\langle\varepsilon\rangle}$ of HML is the language defined inductively on the syntax:

$$\phi \ ::= \ \texttt{true} \mid \langle\varepsilon\rangle\langle a\rangle\phi \mid \phi\wedge\phi \mid \langle\varepsilon\rangle(\phi\wedge\phi)$$

where, for any ϕ_1 and $\phi_2\in\text{HML}_{\text{r}}^{\langle\varepsilon\rangle}$, $\phi_1\wedge\phi_2$ and $\langle\varepsilon\rangle(\phi_1\wedge\phi_2)$ exists in $\text{HML}_{\text{r}}^{\langle\varepsilon\rangle}$ if and only if $\text{act}(\phi_1)\cap\text{act}(\phi_2)=\emptyset$.

Definition 3.5 The sublanguage $\mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$ of $\mathtt{HML}$ is the language defined inductively on the syntax:

$$\psi \ ::= \ \mathtt{true} \mid [\varepsilon]\langle a\rangle\psi \mid \psi \wedge \psi \mid [\varepsilon](\psi \wedge \psi)$$

where, for any ψ_1 and $\psi_2 \in \mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$, $\psi_1 \wedge \psi_2$ and $[\varepsilon](\psi_1 \wedge \psi_2)$ exists in $\mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$ if and only if $\mathrm{act}(\psi_1) \cap \mathrm{act}(\psi_2) = \emptyset$.

It should be clear that the syntactic restrictions we have imposed on the logics ensure that our interpretation is probabilistically sound, since the probabilities corresponding to $[\![\phi_1]\!]E$ and $[\![\phi_2]\!]E$ are independent for any $E \in \mathcal{R}$ and $\phi_1 \wedge \phi_2 \in \mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle} \cup \mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$.

We now establish connections between $\mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle}$ and $\mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$ and our testing preorder, for which the following lemmas and definitions are required. For the remainder of this section we only give the proofs relating to $\mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle}$, as the cases for $\mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$ follow similarly.

Lemma 3.6 *For any $\{\phi_1, \ldots, \phi_m\} \subseteq \mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle}$, if $\mathrm{act}(\phi_i) \cap \mathrm{act}(\phi_j) = \emptyset$ for all $1 \leq i \neq j \leq m$ then there exists $\phi \in \mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle}$ such that*

$$\mathrm{act}(\phi) = \bigcup_{i=1}^{m} \mathrm{act}(\phi_i),\ [\![\phi]\!]E = \left[\!\!\left[\langle\varepsilon\rangle \left(\bigwedge_{i=1}^{m} \phi_i \right) \right]\!\!\right] E \text{ and } [\![\phi]\!]S = \left[\!\!\left[\bigwedge_{i=1}^{m} \phi_i \right]\!\!\right] S$$

for all $E \in \mathcal{R}$ and $S \in \mathcal{P}_{\mathrm{fr}}(\mathcal{A}ct \times \mu(\mathcal{R}))$.

Proof. The proof is by induction on $m \in \mathbb{N}$. □

Lemma 3.7 *For any $\{\psi_1, \ldots, \psi_m\} \subseteq \mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$, if $\mathrm{act}(\psi_i) \cap \mathrm{act}(\psi_j) = \emptyset$ for all $1 \leq i \neq j \leq m$ then there exists $\psi \in \mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$ such that*

$$\mathrm{act}(\psi) = \bigcup_{i=1}^{m} \mathrm{act}(\psi_i),\ [\![\psi]\!]E = \left[\!\!\left[[\varepsilon] \left(\bigwedge_{i=1}^{m} \psi_i \right) \right]\!\!\right] E \text{ and } [\![\psi]\!]S = \left[\!\!\left[\bigwedge_{i=1}^{m} \psi_i \right]\!\!\right] S$$

for all $E \in \mathcal{R}$ and $S \in \mathcal{P}_{\mathrm{fr}}(\mathcal{A}ct \times \mu(\mathcal{R}))$.

Definition 3.8 If $(\!|r|\!) \in \mathtt{T}$, put $r \parallel \bot = \omega \parallel r = r$, and if $(\!|r_1|\!), (\!|r_2|\!) \in \mathtt{T}$ are such that $r_1 = [a_1.T_1, \ldots, a_m.T_m]$, $r_2 = [a'_1.T'_1, \ldots, a'_m.T'_{m'}]$, and r_1 and r_2 are independent, put: $r_1 \parallel r_2 = [a_1.T_1, \ldots, a_m.T_m, a'_1.T'_1, \ldots, a'_m.T'_{m'}]$.

Furthermore, if $T_1 = (t_1, \ldots, t_m) \in \mathtt{T}_\omega$ and $T_2 = (t'_1, \ldots, t'_{m'}) \in \mathtt{T}_\omega$ such that T_1 and T_2 are independent, put: $T_1 \parallel T_2 = (t_1, \ldots, t_m, t'_1, \ldots, t'_{m'})$.

Lemma 3.9 *If $(\!|r_1|\!), (\!|r_2|\!) \in \mathtt{T}$ and $r_1 \parallel r_2$ is defined, then $(\!|r_1 \parallel r_2|\!) \in \mathtt{T}$ and for all $S \in \mathcal{P}_{\mathrm{fr}}(\mathcal{A}ct \times \mu(\mathcal{R}))$: $\mathsf{R}_*(S)(r_1 \parallel r_2) = \mathsf{R}_*(S)(r_1) \cdot \mathsf{R}_*(S)(r_2)$.*

Lemma 3.10 *If $T_1, T_2 \in \mathtt{T}_\omega$ and $T_1 \parallel T_2$ is defined, then $T_1 \parallel T_2 \in \mathtt{T}_\omega$ and for all $E \in \mathcal{R}$: $\mathsf{R}_*(E)(T_1 \parallel T_2) = \mathsf{R}_*(E)(T_1) \cdot \mathsf{R}_*(E)(T_2)$.*

We are now able to state a fundamental connection between our testing scenario and the quantitative interpretation of $\mathtt{HML}$: there is a bijective correspondence between the tests and formulae of $\mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle}$ (respectively $\mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$) such that, for every process, the probabilities assigned to the tests via the map $\mathsf{R}_{\mathbf{glb}}$ (outcomes of random experiments) agree with the probabilities assigned to the formulae of $\mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle}$ (respectively $\mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$) via our interpretation.

Proposition 3.11 *For all $t \in \mathtt{T}$ there exists $\phi_t \in \mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle}$ such that for all $E \in \mathcal{R}$, $[\![\phi_t]\!]E = \mathsf{R}_{\mathbf{lub}}(E)(t)$.*

Proof. The proposition is proved by induction on $t \in \mathtt{T}$.

If $t = (\!|\omega|\!)$, then we set $\phi_{(\!|t|\!)} = \mathtt{true}$, and the proposition holds by definition of $\mathsf{R}_{\mathbf{lub}}$ and $[\![\cdot]\!]$.

If $t = (\!|[a_1.T_1, \ldots, a_m.T_m]|\!)$, then T_i is of the form $(t_1^i, \ldots, t_{m_i}^i)$ for all $1 \leq i \leq m$ and in this case using Lemma 3.6 we set ϕ_t to the formula of $\mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle}$ such that:

$$[\![\phi_t]\!]E = \left[\!\!\left[\langle\varepsilon\rangle \left(\bigwedge_{i=1}^{m} (\langle a_i\rangle \phi_{T_i}) \right) \right]\!\!\right] E$$

for all $E \in \mathcal{R}$, and $\phi_{T_i} = \bigwedge_{j=1}^{m_i} \phi_{t_j^i}$ for all $1 \leq i \leq m$. Now, for any $F \in \mathcal{R}$ and $1 \leq i \leq m$ by definition of $[\![\cdot]\!]$ and ϕ_{T_i}:

$$\begin{aligned}
[\![\phi_{T_i}]\!]F &= \prod_{j=1}^{m_i} [\![\phi_{t_j^i}]\!]F \\
&= \prod_{j=1}^{m_i} \mathsf{R}_{\mathbf{lub}}(F)(t_j^i) \text{ by induction} \\
&= \mathsf{R}_{\mathbf{lub}}(F)(T_i) \qquad \text{by definition of } \mathsf{R}_{\mathbf{lub}}.
\end{aligned}$$

Next, for any $S \in \mathcal{P}_{\mathrm{fr}}(\mathcal{A}ct \times \mu(\mathcal{R}))$ by definition of $[\![\cdot]\!]$:

$$\begin{aligned}
[\![\langle a_i\rangle \phi_{T_i}]\!]S &= \begin{cases} \sum_{F \in \mathcal{R}} \pi(F) \cdot [\![\phi_{T_i}]\!]F \text{ if } (a_i, \pi) \in S \text{ for some } \pi \in \mu(\mathcal{R}) \\ \qquad 0 \qquad\quad \text{otherwise} \end{cases} \\
&= \begin{cases} \sum_{F \in \mathcal{R}} \pi(F) \cdot \mathsf{R}_{\mathbf{lub}}(F)(T_i) \text{ if } (a_i, \pi) \in S \text{ for some } \pi \in \mu(\mathcal{R}) \\ \qquad 0 \qquad\quad \text{otherwise} \end{cases} \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{from above} \\
&= \mathsf{R}_{\mathbf{lub}}(S)(a_i.T_i) \qquad\qquad\qquad\qquad\qquad\qquad \text{by definition of } \mathsf{R}_{\mathbf{lub}}.
\end{aligned}$$

Finally, by construction of ϕ_t for any $E \in \mathcal{R}$:

$$\begin{aligned}
[\![\phi_t]\!]E &= \left[\!\!\left[\langle\varepsilon\rangle \left(\bigwedge_{i=1}^{m} (\langle a_i\rangle \phi_{T_i}) \right) \right]\!\!\right] E \\
&= \max_{E \to S} \prod_{i=1}^{m} [\![\langle a_i\rangle \phi_{T_i}]\!]S \qquad\qquad \text{by definition of } [\![\cdot]\!] \\
&= \max_{E \to S} \prod_{i=1}^{m} \mathsf{R}_{\mathbf{lub}}(S)(a_i.T_i) \qquad\qquad \text{from above} \\
&= \max_{E \to S} \mathsf{R}_{\mathbf{lub}}(S)([a_1.T_1, \ldots, a_m.T_m]) \text{ by definition of } \mathsf{R}_{\mathbf{lub}} \\
&= \mathsf{R}_{\mathbf{lub}}(E)(t) \qquad\qquad \text{by definition of } \mathsf{R}_{\mathbf{lub}}
\end{aligned}$$

as required. □

Proposition 3.12 *For all $t \in \mathtt{T}$ there exists $\psi_t \in \mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$ such that for all $E \in \mathcal{R}$, $[\![\psi_t]\!]E = \mathsf{R}_{\mathbf{glb}}(E)(t)$.*

Proposition 3.13 *For all $\phi \in \mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle}$ there exists $(\!|r_\phi|\!) \in \mathtt{T}$ and $T_\phi \in \mathtt{T}_\omega$ such that $[\![\phi]\!]S = \mathsf{R}_{\mathbf{lub}}(S)(r_\phi)$ for all $S \in \mathcal{P}_{\mathrm{fr}}(\mathcal{A}ct \times \mu(\mathcal{R}))$, and $[\![\phi]\!]E = \mathsf{R}_{\mathbf{lub}}(E)(T_\phi)$ for all $E \in \mathcal{R}$.*

Proof. The proof follows by induction on $\phi \in \mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle}$, by putting:

$$r_\phi = \begin{cases} \omega & \text{if } \phi = \mathtt{true} \\ [a.T_{\phi'}] & \text{if } \phi = \langle\varepsilon\rangle\langle a\rangle\phi' \\ r_{\phi_1} \parallel r_{\phi_2} & \text{if } \phi = \phi_1 \wedge \phi_2 \\ r_{\phi_1} \parallel r_{\phi_2} & \text{if } \phi = \langle\varepsilon\rangle(\phi_1 \wedge \phi_2) \end{cases}$$

and

$$T_\phi = \begin{cases} (\!|\omega|\!) & \text{if } \phi = \mathtt{true} \\ (\!|[a.T_{\phi'}]|\!) & \text{if } \phi = \langle\varepsilon\rangle\langle a\rangle\phi' \\ T_{\phi_1} \parallel T_{\phi_2} & \text{if } \phi = \phi_1 \wedge \phi_2 \\ (\!|r_{\phi_1} \parallel r_{\phi_2}|\!) & \text{if } \phi = \langle\varepsilon\rangle(\phi_1 \wedge \phi_2). \end{cases}$$

□

Proposition 3.14 *For all $\psi \in \mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$ there exists $(\!|r_\psi|\!) \in \mathtt{T}$ and $T_\psi \in \mathtt{T}_\omega$ such that $[\![\psi]\!]S = \mathsf{R}_{\mathbf{glb}}(S)(r_\psi)$ for all $S \in \mathcal{P}_{\mathrm{fr}}(\mathcal{A}ct \times \mu(\mathcal{R}))$, and $[\![\psi]\!]E = \mathsf{R}_{\mathbf{glb}}(E)(T_\psi)$ for all $E \in \mathcal{R}$.*

Finally, using Proposition 3.11, Proposition 3.12, Proposition 3.13 and Proposition 3.14, we prove the central theorem connecting $\mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle} \cup \mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$ and $\sqsubseteq^{\mathsf{R}}$ missing from [14].

Theorem 3.15 *For all E, $F \in \mathcal{R}$, $E \sqsubseteq^{\mathsf{R}} F$ if and only if $[\![\phi]\!]E \leq [\![\phi]\!]F$ for all $\phi \in \mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle}$ and $[\![\psi]\!]E \leq [\![\psi]\!]F$ for all $\psi \in \mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$.*

We point out that adding negation to the quantitative HML is rather delicate, and refer the interested reader to [15,16]. However, if we restrict ourselves to deterministic probabilistic transition systems then adding negation to the logic HML does *not* influence the equivalence induced from the logic, in the sense that thus obtained equivalence will still correspond to the restriction of the equivalence $\overset{\mathsf{R}}{\sim}$ to the deterministic probabilistic transition systems [23].

4 Adding Fixed Point Operators to HML

In this section we add a fixed point operator to the logics $\mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle}$ and $\mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$ and compare the results with our maps $\mathsf{R}_{\mathbf{lub}}$ and $\mathsf{R}_{\mathbf{glb}}$ respectively. This provides the probabilistic justification for the quantitative interpretation of the modal mu-calculus missing from [14]. We note that we only prove results relating to $\mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle}$ and $\mathsf{R}_{\mathbf{lub}}$, as the results for $\mathtt{HML}_{\mathrm{r}}^{[\varepsilon]}$ and $\mathsf{R}_{\mathbf{glb}}$ are dual.

To add a fixed point operator to our logic we must first add variables (ranged over by Var) to the syntax of $\mathtt{HML}_{\mathrm{r}}^{\langle\varepsilon\rangle}$ to form $\mathtt{HML}_{\mathrm{rv}}^{\langle\varepsilon\rangle}$. As usual we then

extend the maps $[\![\cdot]\!]$ and $\mathsf{R}_{\mathbf{lub}}$ by means of environments $\rho : \mathtt{Var} \longrightarrow (\mathcal{R} \longrightarrow [0,1])$ so that $[\![\phi]\!]\rho : \mathcal{R} \longrightarrow [0,1]$ for any $\phi \in \mathtt{HML}^{\langle\varepsilon\rangle}_{\mathrm{rv}}$, and likewise for $\mathsf{R}_{\mathbf{lub}}$. We omit the environments to simplify the notation. To compare the tests of $\mathtt{T}$ to fixed point operators of $\mathtt{HML}^{\langle\varepsilon\rangle}_{\mathrm{rv}}$, we construct *unfoldings* of formulae, and using the map between formulae of $\mathtt{HML}^{\langle\varepsilon\rangle}_{\mathrm{rv}}$ and $\mathtt{T}$ given in Proposition 3.13 we then consider these unfoldings as elements of our testing language. Formally, we have the following definitions.

Definition 4.1 For all $\phi \in \mathtt{HML}^{\langle\varepsilon\rangle}_{\mathrm{rv}}$ and $x \in \mathtt{Var}$, we define ϕ^n_x by induction on $n \in \mathbb{N}$ as follows: $\phi^0_x = \mathtt{true}$ and $\phi^{n+1}_x = \phi\{\phi^n_x/x\}$.

Definition 4.2 For any $\phi \in \mathtt{HML}^{\langle\varepsilon\rangle}_{\mathrm{rv}}$, let $(\!|r_\phi|\!) \in \mathtt{T}$ and $T_\phi \in \mathtt{T}_\omega$ be the tests defined by induction on $\phi \in \mathtt{HML}^{\langle\varepsilon\rangle}_{\mathrm{rv}}$ as follows:

$$r_\phi = \begin{cases} \phi & \text{if } \phi \in \mathtt{Var} \\ \omega & \text{if } \phi = \mathtt{true} \\ [a.T_{\phi'}] & \text{if } \phi = \langle\varepsilon\rangle\langle a\rangle\phi' \\ r_{\phi_1} \parallel r_{\phi_2} & \text{if } \phi = \phi_1 \wedge \phi_2 \text{ or } \phi = \langle\varepsilon\rangle(\phi_1 \wedge \phi_2) \end{cases}$$

$$\text{and} \quad T_\phi = \begin{cases} \phi & \text{if } \phi \in \mathtt{Var} \\ (\!|\omega|\!) & \text{if } \phi = \mathtt{true} \\ (\!|[a.T_{\phi'}]|\!) & \text{if } \phi = \langle\varepsilon\rangle\langle a\rangle\phi' \\ T_{\phi_1} \parallel T_{\phi_2} & \text{if } \phi = \phi_1 \wedge \phi_2 \\ (\!|r_{\phi_1} \parallel r_{\phi_2}|\!) & \text{if } \phi = \langle\varepsilon\rangle(\phi_1 \wedge \phi_2) \end{cases}$$

where for any $S \in \mathcal{P}_{\mathrm{fr}}(\mathcal{A}ct \times \mu(\mathcal{R}))$, $E \in \mathcal{R}$ and $x \in \mathtt{Var}$ let $\mathsf{R}_{\mathbf{lub}}(S)(x) = \rho(S)$ and $\mathsf{R}_{\mathbf{lub}}(E)(x) = \rho(E)$.

Then, similarly to Proposition 3.13, we have the following proposition.

Proposition 4.3 *For all $\phi \in \mathtt{HML}^{\langle\varepsilon\rangle}_{\mathrm{rv}}$, there exists $(\!|r_\phi|\!) \in \mathtt{T}$ and $T_\phi \in \mathtt{T}_\omega$ such that $[\![\phi]\!]\rho S = \mathsf{R}_{\mathbf{lub}}(S)(r_\phi)$ for all $S \in \mathcal{P}_{\mathrm{fr}}(\mathcal{A}ct \times \mu(\mathcal{R}))$, and $[\![\phi]\!]\rho E = \mathsf{R}_{\mathbf{lub}}(E)(T_\phi)$ for all $E \in \mathcal{R}$.*

The sequence of formulae given in Definition 4.1 (fixed point unfoldings) gives rise, via the map between formulae and tests given in Definition 4.2, to the sequences of tests $\langle r^n_{\phi_x}\rangle_{n\in\mathbb{N}}$, $\langle T^n_{\phi_x}\rangle_{n\in\mathbb{N}}$. The following lemma and proposition demonstrate that successive unfoldings *improve* the probability upper bound obtained with the help of the map $\mathsf{R}_{\mathbf{lub}}$.

Lemma 4.4 *If $\phi, \theta_1, \theta_2 \in \mathtt{HML}^{\langle\varepsilon\rangle}_{\mathrm{rv}}$ and $x \in \mathtt{Var}$ such that*

$$\mathsf{R}_{\mathbf{lub}}(S)(r_{\theta_1}) \leq \mathsf{R}_{\mathbf{lub}}(S)(r_{\theta_2}) \text{ and } \mathsf{R}_{\mathbf{lub}}(E)(T_{\theta_1}) \leq \mathsf{R}_{\mathbf{lub}}(E)(T_{\theta_2})$$

for all $S \in \mathcal{P}_{\rm fr}(\mathcal{A}ct \times \mu(\mathcal{R}))$ *and* $E \in \mathcal{R}$*, then*

$$(i)\ \mathsf{R}_{\mathbf{lub}}(S)(r_{\phi\{\theta_1/x\}}) \leq \mathsf{R}_{\mathbf{lub}}(S)(r_{\phi\{\theta_2/x\}}) \quad \textit{for all } S \in \mathcal{P}_{\rm fr}(\mathcal{A}ct \times \mu(\mathcal{R}))$$

$$(ii)\ \mathsf{R}_{\mathbf{lub}}(E)(T_{\phi\{\theta_1/x\}}) \leq \mathsf{R}_{\mathbf{lub}}(E)(T_{\phi\{\theta_2/x\}}) \textit{ for all } E \in \mathcal{R}.$$

Proof. The proof follows by induction on $\phi \in \mathtt{HML}^{\langle\varepsilon\rangle}_{\rm rv}$. □

Proposition 4.5 *For all* $S \in \mathcal{P}_{\rm fr}(\mathcal{A}ct \times \mu(\mathcal{R}))$*,* $E \in \mathcal{R}$*,* $\phi, \theta \in \mathtt{HML}^{\langle\varepsilon\rangle}_{\rm rv}$ *and* $x \in \mathtt{Var}$*:* $\mathsf{R}_{\mathbf{lub}}(S)(r_{\theta\{\phi^{n+1}_x/x\}}) \leq \mathsf{R}_{\mathbf{lub}}(S)(r_{\theta\{\phi^n_x/x\}})$ *and* $\mathsf{R}_{\mathbf{lub}}(E)(T_{\theta\{\phi^{n+1}_x/x\}}) \leq \mathsf{R}_{\mathbf{lub}}(E)(T_{\theta\{\phi^n_x/x\}})$*.*

Proof. The proof is by induction on $\theta \in \mathtt{HML}^{\langle\varepsilon\rangle}_{\rm rv}$. □

Corollary 4.6 *For all* $\phi \in \mathtt{HML}^{\langle\varepsilon\rangle}_{\rm rv}$*,* $x \in \mathtt{Var}$ *and* $E \in \mathcal{R}$*, the limit*

$$\lim_{n\to\infty} \mathsf{R}_{\mathbf{lub}}(E)(T^n_{\phi_x})$$

exists and is in the interval $[0, 1]$*.*

Proof. If we consider any $\phi \in \mathtt{HML}^{\langle\varepsilon\rangle}_{\rm rv}$, then using Proposition 4.5 and letting $\theta = \phi$ we have $\langle \mathsf{R}_{\mathbf{lub}}(E)(T^n_{\phi_x})\rangle_{n\in\mathbb{N}}$ is a decreasing sequence in the interval $[0, 1]$, and hence the (unique) limit exists and is in the interval $[0, 1]$. □

From Corollary 4.6 we know that the limit of probability upper bounds for successive unfoldings exists. This limit $\lim_{n\to\infty} \mathsf{R}_{\mathbf{lub}}(E)(T^n_{\phi_x})$ is in fact the value of the greatest fixed point $\nu x.\phi$ in Huth and Kwiatkowska's interpretation of the modal mu-calculus (which does not deal with non-determinism). For any $\phi \in \mathtt{HML}^{\langle\varepsilon\rangle}_{\rm rv}$ and $E \in \mathcal{R}$ we have:

$$[\![\nu x.\phi]\!]E = \lim_{n\to\infty} \mathsf{R}_{\mathbf{lub}}(E)(T^n_{\phi_x}).$$

The connection with the greatest, as opposed to the least, fixed point operator arises from the fact that there is no test representing `false` in our testing language `T`, and hence we must begin all iterations from `true` (that is, $(\!|\omega|\!)$), and since $\mathsf{R}_{\mathbf{lub}}(E)(T) \leq 1$ for all $E \in \mathcal{R}$ and $T \in \mathtt{T}_\omega$, any monotone sequence we construct will either be constant at 1 or *decreasing*. Hence, the limit corresponds with the greatest fixed point. In general, the values of the greatest fixed point operator with respect to the formulae of $\mathtt{HML}^{\langle\varepsilon\rangle}_{\rm rv}$ and $\mathtt{HML}^{[\varepsilon]}_{\rm r}$ may differ if processes are non-deterministic. Intuitively, the pair of values

$$[[\![\nu x.[\varepsilon]\langle a\rangle]\!]E, [\![\nu x.\langle\varepsilon\rangle\langle a\rangle]\!]E]$$

corresponds to the *interval* containing the probability that $\tilde{E}$ will perform an infinite path of a actions.

5 Conclusions

We have formulated a testing equivalence on reactive probabilistic processes which exhibit three kinds of choice: action-guarded probabilistic, external and internal. The equivalence is non-probabilistically branching time, but, unlike

probabilistic bisimulation, does not make distinctions according to when the probabilistic choices are made. Such situations arise when actions executed by the process have no effect on random choices, for example when selecting an option from a menu will not influence the outcome of a coin toss.

The derived equivalence is a congruence for a subcalculus of CSP including internal, external and probabilistic choice as well as synchronous parallel (for details see [18,23], where also fully abstract denotational semantics is presented), but not for hiding and asynchronous parallel. Asynchronous parallel is important in the compositional specification and verification of probabilistic protocols consisting of independently acting components. Hiding is relied upon when verifying CSP processes against specifications by means of `fdr2`. In such situations, and also when probabilities are affected by external actions, our equivalence is not appropriate. A potential solution to this problem, proposed in [23], is to first consider a transition system model where processes perform internal probabilistic choices of the kind $E \sqcap_p F$, where $p + q = 1$, meaning $E \sqcap_p F$ will act as the process E with probability p and F with probability q, as opposed to action-guarded internal probabilistic choices as presented here. The hope is that such a model would admit the full calculus of CSP [7] extended with an internal probabilistic choice operator without losing the congruence property of the equivalence.

We have also given a logical characterization of the equivalence in terms of the quantitative version of `HML` of [14] extended with silent actions, and have established its probabilistic soundness not dealt with in [14], albeit under syntactic restrictions that impose independence. These could be removed at a cost of introducing conditional probabilities. The parallels between our approach and that of [5], particularly as far as the computing the least upper bounds (and their duals) on probabilities are concerned, mean that the model checking algorithm of [5] should apply to our case also. Adding negation to our setting is more difficult since negating a lower bound yields an upper bound, and vice versa [16]; a suitable framework is proposed in [15].

Acknowledgement

We are grateful to Michael Huth, Achim Jung, Gavin Lowe and Bill Roscoe for comments on earlier versions of this work.

References

[1] S. Abramsky. Observational equivalence as a testing equivalence. *Theoretical Computer Science, 53:225-241, 1987.*

[2] C. Baier. Polynomial time algorithms for testing probabilistic bisimulation and simulation. *In Proc. CAV'96, volume 1102 of Lecture Notes in Computer Science, pages 38-49, Springer Verlag, 1996.*

[3] C. Baier and M.Z. Kwiatkowska. Domain equations for probabilistic processes (Extended Abstract). *In Proc. EXPRESS Workshop, volume 7, Electronic Notes in Theoretical Computer Science, Elsevier, 1997.*

[4] J.A. Bergstra and J.W. Klop. Process algebra for synchronous communication. *Information and Computation, 60:109-134, 1984.*

[5] A. Bianco and L. de Alfaro. Model checking of probabilistic and nondeterministic systems. *Proc. Foundations of Software Technology and Theoretical Computer Science, in volume 1026 of Lecture Notes in Computer Science, pages 499-513, Springer Verlag, 1995.*

[6] B. Bloom and A.R. Meyer. A remark on bisimulation between probabilistic processes. *In A.R. Meyer and M.A. Taitslin, editors, Symp. on Logical Foundations of Computer Science, volume 363 of Lecture Notes in Computer Science, pages 26-40, Springer Verlag, 1989.*

[7] S.D. Brookes, C.A.R. Hoare and A.W. Roscoe. A theory of communicating sequential processes. *Journal of the ACM, 31(3):560-599, 1984.*

[8] J. Desharnais, A. Edalat and P. Panangaden. A logical characterization of bisimulation for labeled Markov processes. *In Proc. 13th IEEE Symposium on Logic in Computer Science (LICS), pages 478-487, 1998.*

[9] A. Giacalone, C-C. Jou and S.A. Smolka. Algebraic reasoning for probabilistic concurrent systems. *In M. Broy and C.B. Jones, editors, Proc. IFIP TC2 Working Conference on Programming Concepts and Methods, Sea of Galilee, Israel, pages 443-458, 1990.*

[10] R.J. van Glabbeek. The linear time-branching time spectrum. *In J.C.M. Beaten and J.W. Klop editors, Proc. CONCUR'90, volume 458 of Lecture Notes in Computer Science, pages 278-297, Springer Verlag, 1990.*

[11] R.J. van Glabbeek, S.A. Smolka, B. Steffen and C.M.N. Tofts. Reactive, generative and stratified models of probabilistic processes. *In Proc. 5th IEEE Int. Symp. on Logic in Computer Science (LICS), pages 130-141, 1990.*

[12] H.A. Hansson. Time and probability in the formal design of distributed systems. *Volume 1 of Real-Time Safety Critical Systems, Elsevier, 1994.*

[13] M.C.B. Hennessy and R. Milner. Algebraic laws for nondeterminism and concurrency. *Journal of the Association for Computing Machinery, 32(1):137-161, 1985.*

[14] M. Huth and M.Z. Kwiatkowska. Quantitative analysis and model checking. *In Proc. 12th IEEE Int. Symp. on Logic in Computer Science (LICS), pages 111-122, 1997.*

[15] M. Huth. The interval domain: a matchmaker for aCTL and aPCTL. *Technical Report CIS-97-17, Department of Computing and Information Sciences, Kansas State University, 1997.*

[16] M. Huth and M.Z. Kwiatkowska. Comparing CTL and PCTL on labeled Markov chains. *In PROCOMET'98, IFIP, Chapman & Hall, 1998.*

[17] C-C. Jou and S.A. Smolka. Equivalences, congruences and complete axiomatisations for probabilistic processes. *In J.C.M. Baeten and J.W. Klop, editors, CONCUR'90, volume 458 of Lecture Notes in Computer Science, pages 367-383, Springer Verlag, 1990.*

[18] M.Z. Kwiatkowska and G.J. Norman. A fully abstract metric-space denotational semantics for reactive probabilistic processes. *In Proc. COMPROX'98, volume 13 of Electronic Notes in Theoretical Computer Science, Elsevier, 1998.*

[19] K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation, 94(1):1-28, 1991.*

[20] G. Lowe. Representing nondeterministic and probabilistic behaviour in reactive processes. *Technical Report, 1993.* Available at `http://www.mcs.le.ac.uk/~glowe/Publications.html`.

[21] R. Milner. Communication and concurrency. *Prentice Hall, 1989.*

[22] C. Morgan, A. McIver, K. Seidel and J.W. Sanders. Refinement-oriented probability for CSP. *Formal Aspects of Computing, 8(6):617-647, 1996.*

[23] G.J. Norman. Metric semantics for reactive probabilistic processes. *Ph.D Thesis, School of Computer Science, The University of Birmingham, November 1997.*

[24] A.W. Roscoe. The theory and practice of concurrency. *Prentice Hall International Series in Computer Science, 1997.*

[25] R. Segala and N. Lynch. Probabilistic simulations for probabilistic processes. *In B. Jonsson and J. Parrow, editors, Proc. CONCUR'94, volume 836 of Lecture Notes in Computer Science, pages 481-496, Springer, 1994.*

[26] K. Seidel. Probabilistic communicating processes. *Theoretical Computer Science, 152:219-249, 1995.*

[27] Wang Yi and K.G. Larsen. Testing probabilistic and non-deterministic processes. *Protocol Specification, Testing and Verification XII:47-61, Florida, USA, 1992.*