Acta Informatica
https://doi.org/10.1007/s00236-020-00379-2

ORIGINAL ARTICLE

On the probabilistic bisimulation spectrum with silent moves

Christel Baier[1] · Pedro R. D'Argenio[2,3,4] · Holger Hermanns[4,5]

Received: 9 February 2020 / Accepted: 26 February 2020

Abstract
In this paper we look at one of the seminal works of Rob van Glabbeek from a probabilistic angle. We develop the bisimulation spectrum with silent moves for probabilistic models, namely Markov decision processes. Especially the treatment of divergence makes this endeavour challenging. We provide operational as well as logical characterisations of a total of 32 bisimilarities.

1 Introduction

The probably most resonating contribution brought to computer science through the conception of concurrency theory is manifested by behavioural relations. Especially bisimulation relations are the prime vehicle to equate, respectively distinguish processes according to the behaviour they exhibit when interacting with other processes, taking the stepwise behaviour of states in labeled transition systems as a reference.

The systematic and comparative study of such relations has been pioneered by Rob van Glabbeek in a pair of landmark papers providing a concise yet precise overview of the linear time—branching time spectrum [23,24]. Arguably, bisimulation relations form the core of

This work has received financial support by DFG Grant 389792660 as part of TRR 248 (see perspicuous-computing.science), by ERC Advanced Grant 69561 (POWVER), by ANPCyT PICT-2017-3894 (RAFTSys), and by SeCyT-UNC 33620180100354CB (ARES), by the Cluster of Excellence EXC 2050/1 (CeTI, Project ID 390696704, as part of Germany's Excellence Strategy), by the DFG-Projects BA-1679/11-1 and 1679/12-1 and by the CRC 912 HAEC.

✉ Christel Baier
christel.baier@tu-dresden.de

Pedro R. D'Argenio
dargenio@famaf.unc.edu.ar

Holger Hermanns
hermanns@cs.uni-saarland.de

1 Institute of Theoretical Computer Science, Technische Universität Dresden, Dresden, Germany

2 FAMAF, Universidad Nacional de Córdoba, Buenos Aires, Argentina

3 CONICET, Buenos Aires, Argentina

4 Saarland University, Saarland Informatics Campus, Saarbrücken, Germany

5 Institute of Intelligent Software, Guangzhou, China

Published online: 06 May 2020

the finer branching time fragment of this spectrum. A bisimulation equates processes whose behaviours (in terms of moves from state to state) can be mutually simulated [18].

In strong bisimulation, each individual move needs to be mimicked. But often in concurrency theory one faces internal computations that are to be abstracted away by a natural notion of behavioural equivalence. These are “silent” moves in the jargon of van Glabbeek, opposed to observable moves. A variety of bisimulations for labeled transition systems can be defined capturing the essence of observable behaviour in varying degrees.

The finest relation in this bisimulation spectrum with silent moves is branching bisimulation [26] with explicit divergence [2,23], the coarsest is Milner’s notion of weak bisimulation [17], which is blind wrt. divergence. The *bisimulation spectrum with silent moves* spans between these two extremes. Explicit divergence has lately been refined [16] with respect to the (in)ability to step through distinguishable behaviour while diverging.

In this contribution, we navigate through the *probabilistic bisimulation spectrum with silent moves*. Simply speaking, probabilities enter the classical spectrum by making the effect of moves probabilistic. Instead of representing moves from state to state, transitions represent moves from states to probability distributions over states, yielding the model world of Markov decision processes (MDP), or better, probabilistic automata [20]. The classical setting then becomes a special case, obtained if only allowing distributions with a singleton support set (termed “Dirac” distributions.)

Silent moves are easy to capture in the classical setting (basically as the transitive closure of internal transitions), while this gets technically more involved in the MDP setting [8,10,20]. We work with combined and infinitary silent moves, even though we slighty depart from the traditional approaches for the sake of elegance. We review the meaningful generalisations of the known divergence and bisimulation concepts, encompassing operational and modal logic characterisations of a total of 32 bisimulation relations. In doing so, we attempt to work in the style of Rob van Glabbeek wrt. providing crisp, precise, and modular definitions. Grosso modo, we find that the classical spectrum extends smoothly with some refinements being needed to properly incorporate probabilistic divergence, all the while the necessary proofs partially become very intricate. Furthermore, we need to restrict to finite-state models for some of the more advanced notions of explicit divergence, namely almost-sure explicit divergence, and explicit divergence with positive probability.

Related work Apart from the seminal papers by Rob van Glabbeek [23,24] introducing the basic considerations on which our work rests, there are several subsequent works that have developed aspects of the non-probabilistic bisimulation spectrum with silent moves. Explicit divergence in branching bisimulation has been carefully studied together with modal logic characterisations [25], and the notion of explicit divergence has been further refined lately [16]. In the probabilistic setting, modal logics for strong bisimulation and simulation relations have been the topic of [14] and were discussed for strong relations in continuous spaces [5] as well as on distributions [13].

Probabilistic bisimulation relations that abstract from internal computations have been studied from a logical perspective, too, albeit generally without any specific consideration of divergence phenomena—contrary to the non-probabilistic setting. Concretely, the temporal logics PCTL and PCTL* have been studied in [9] for labeled concurrent Markov chains (i.e., a subclass of probabilistic automata where the silent moves are purely probabilistic), and have been matched by incrementally defined notions of strong bisimulations in the state-based (rather than action-labeled) setting [21]. Matching notions of bisimulations for the equivalences induced by the next-free fragments of PCTL and PCTL* have been provided in [21] in the style of (state-based) branching bisimulations that depart from standard bisimulation notions by imposing global conditions on the probabilities of stutter-invariant sets of paths so that divergence insensitivity is a built-in feature. Apart from this, distribution-based bisimulations with silent moves have gained popularity [8,10] and subsequently been studied from a logical perspective [15]. The axiomatic perspective on weak bisimulation in the presence of recursion has recently been developed by Fischer and van Glabbeek [11]. So, all in all, there are some few isolated spots in the probabilistic bisimulation spectrum with silent moves that have been shed light on, but a deep and profound discussion of the entire spectrum has so far not been developed. This is what the present paper provides.

Outline We start off in Sect. 2 by defining the essential concepts and notations, most notably our notion of compound transitions. Section 3 then introduces the operational definitions of the bisimulations and divergence notions that span the spectrum. Basic properties of these relations are established in Sect. 4 and in Sect. 5. The latter focusses on the properties of the divergence-sensitive bisimulations and divergence probabilities. It is here where we need to restrict to finite-state behaviours. Section 6 introduces a spectrum of modal logics, which in Sect. 7 is shown to match the operational definitions introduced before, in a (pairwise) sound and complete manner. Section 8 summarises our findings by presenting the entire probabilistic bisimulation spectrum with silent moves. We conclude the paper in Sect. 9 with a brief summary and discussion.

2 Preliminaries

Basic definitions Given a set S, its powerset is 2^S. A (discrete) probability distribution over S is a function $\mu \in S \to [0,1]$ such that its support $\mathrm{supp}(\mu) \overset{\mathrm{def}}{=} \{s \in S \mid \mu(s) > 0\}$ is countable and $\sum_{s \in \mathrm{supp}(\mu)} \mu(s) = 1$. If $\mathrm{supp}(\mu)$ is a singleton, then we call μ a Dirac distribution, and if a Dirac distribution has support $\{s\}$ we often denote the distribution as δ_s. $\mathrm{Dist}(S)$ is the set of all probability distributions over S. If $\mu \in \mathrm{Dist}(S)$ and $T \subseteq S$ then we often write $\mu(T)$ for $\sum_{t \in T} \mu(t)$.

Let R_1 and R_2 be binary relations on S. Then, R_1 is called coarser than R_2 iff $R_2 \subseteq R_1$. and R_1 is called strictly coarser than R_2 iff $R_2 \subsetneq R_1$. The notions "finer" and "strictly finer" have analogous meanings.

If R is an equivalence relation on S then we write S/R for the quotient space, i.e., the set of R-equivalence classes. The lifting R to an equivalence $\equiv_R$ on $\mathrm{Dist}(S)$ defined by:

$$\mu \equiv_R \nu \quad \text{iff} \quad \mu(C) = \nu(C) \quad \text{for all } C \in S/R$$

It is easy to see that $\equiv_R$ is indeed an equivalence. We often shall use the fact that if $\mu \in \mathrm{Dist}(S)$ and $s \in S$ then:

$$\mu \equiv_R \delta_s \ \text{ iff } \ \mathrm{supp}(\mu) \subseteq [s]_R$$

Another property of $\equiv_R$ that shall be used at multiple places is the following observation. If $(\mu_i)_{i \in I}$ is a countable family of distributions with $\mu_i \equiv_R \nu$ and $p_i \in]0,1[$ for $i \in I$ such that $\sum_{i \in I} p_i = 1$ then $\sum_{i \in I} p_i \cdot \mu_i \equiv_R \nu$.

Markov decision processes We consider a finite action set Act containing the distinguished *internal* action τ. All other elements of Act are *visible* (or *observable*). We restrict to finite Markov decision processes (MDPs), defined as follows.

Definition 1 An MDP $\mathcal{M}$ is a tuple $(S, \text{Act}, \rightarrow)$ with finite state space S, finite action set Act and finite transition relation

$$\rightarrow \subseteq S \times \text{Act} \times \text{Dist}(S).$$

In what follows, we introduce the notations that are relevant for the other sections. Further details about MDPs can be found, e.g., in [1,19].

A state s of an MDP $\mathcal{M}$ is called *terminal* if there is no outgoing transition of s, i.e., $\{(\alpha, \mu) \in \text{Act} \times \text{Dist}(S) : s \xrightarrow{\alpha} \mu\} = \varnothing$. *Paths* in an MDP are (finite or infinite) sequences

$$\pi = s_0\, \alpha_0\, \mu_0\, s_1\, \alpha_1 \ldots \alpha_{n-1}\, \mu_{n-1}\, s_n \in (S \times \text{Act} \times \text{Dist}(S))^* S \text{ or}$$
$$\pi = s_0\, \alpha_0\, \mu_0\, s_1\, \alpha_1\, \mu_1\, s_2\, \alpha_2\, \mu_2\ \ldots \in (S \times \text{Act} \times \text{Dist}(S))^\omega$$

such that, for each position $i \in \{0, \ldots, n-1\}$ resp. $i \in \mathbb{N}$, $\mathcal{M}$'s transition relation contains $s_i \xrightarrow{\alpha_i} \mu_i$ and $s_{i+1} \in \text{supp}(\mu_i)$. A path is said to be *maximal* if it is either infinite or finite and ends in a terminal state. We write *FinPaths*, *InfPaths*, *MaxPaths* for the sets of finite, infinite resp. maximal paths in $\mathcal{M}$. At various places in the paper, the names of the distributions in paths are irrelevant. In these cases, we will write paths as alternating sequences of states and actions.

A *partial (randomized) scheduler* for $\mathcal{M}$ is a function

$$\sigma : \mathit{FinPaths} \rightarrow \text{Dist}(\rightarrow \cup \{\mathit{stop}\})$$

such that for each finite path $\pi = s_0\, \alpha_0\ \ldots \alpha_{n-1}\, s_n$ and each transition $s \xrightarrow{\alpha} \mu$ in the support of $\sigma(\pi)$ we have $s = s_n$. In particular, if π is maximal (if π's last state s_n is terminal), then σ schedules for π the element *stop* with probability 1, i.e., $\sigma(\pi) = \delta_{\mathit{stop}}$. A partial scheduler σ is called *total*, or briefly a *scheduler*, if

$$\sigma(\pi)(\mathit{stop}) > 0 \text{ iff } \pi \text{ is maximal}$$

A path π is called a σ-path if π can be generated by following σ's decisions. For example, if $\pi = s_0\, \alpha_0\, \mu_0\, s_1\, \alpha_1\, \mu_1\ \ldots$ is an infinite path then π is a σ-path if for each position $i \in \mathbb{N}$, σ's decision for the prefix $s_0\, \alpha_0\, \mu_0\ \ldots \alpha_{i-1}\, \mu_{i-1}\, s_i$ is a distribution Θ where $\Theta(s_i \xrightarrow{\alpha_i} \mu_i) > 0$.

A partial scheduler σ is said to be

- *memoryless* if for all finite paths π_1, π_2 that end in the same state we have: $\sigma(\pi_1) = \sigma(\pi_2)$. In this case, σ can be represented by a function assigning a distribution over transitions to every non-terminal state in $\mathcal{M}$ and assigning *stop* to every terminal state.
- *deterministic* if for each path π, the distribution $\sigma(\pi)$ is a Dirac distribution, in which case σ can be viewed as a partial function from finite paths to transitions.

Given a state s of $\mathcal{M}$ and a (total) scheduler σ, then the behavior of $\mathcal{M}$ under σ can be formalized by an infinite-state tree-like Markov chain $\mathcal{C}_s^\sigma$ where the states are the finite σ-paths starting in state s and with the following transition probability function $P_{\mathcal{C}}$. Let $\pi = s_0\, \alpha_0\, \mu_0 \ldots \alpha_{n-1}\, \mu_{n-1}\, s_n$ be a σ-path from $s = s_0$. If s_n is terminal then so is π as a state of $\mathcal{C}_s^\sigma$. Otherwise $\sigma(\pi)$ is a distribution over $\mathcal{M}$'s transitions from state $t = s_n$. So, there are transitions $t \xrightarrow{\beta_i} \nu_i$ in $\mathcal{M}$ and real numbers $p_1, \ldots, p_k \in\,]0, 1]$ such that $p_1 + \cdots + p_k = 1$ and σ schedules $t \xrightarrow{\beta_i} \nu_i$ for the input path π with probability p_i. Then, the transition probabilities $P_{\mathcal{C}}(\ldots)$ in $\mathcal{C} = \mathcal{C}_s^\sigma$ from π viewed as a state of $\mathcal{C}$ are defined as follows:

$$P_{\mathcal{C}}(\pi, \pi\, \beta_i\, \nu_i\, u) = p_i \cdot \nu_i(u)$$

for the states $u \in \mathrm{supp}(\nu_i)$ and $P_{\mathcal{C}}(\pi, \pi') = 0$ for all other paths π'. So, the successors of π in $\mathcal{C}_s^\sigma$ are the σ-paths π' that extend π by a single step that conforms to σ's decisions. Using standard measure-theoretic concepts, a probability measure $\Pr_s^\sigma(\ldots)$ for measurable sets of maximal paths starting in s is obtained by transfering the standard probability measure of the Markov chain $\mathcal{C}_s^\sigma$ to (measurable) sets of maximal σ-paths in $\mathcal{M}$.

Compound transitions The roots of the MDP model within concurrency theory are probabilistic automata [20], coined by Roberto Segala. One of their basic concepts [8,10,20] is that of *combined transitions* which are obtained as convex combinations of MDP transitions emanating a state. We take the liberty to slightly deviate from the traditional approach, for the sake of easiness of concepts and succinctness of proofs. We instead define *compound* transitions (denoted $\longrightarrow_c$), which for visible actions are convex combinations of ordinary transitions, while compound τ-transition admit non-zero probability for not taking a transition. Formally, for $s \in S$, $\mu \in \mathrm{Dist}(S)$ and $a \in \mathrm{Act}\backslash\{\tau\}$,

- $s \xrightarrow{a}_c \mu$ if there exist transitions $s \xrightarrow{a} \mu_i$ for $i = 1, \ldots, k$, and real numbers $p_1, \ldots, p_k \in]0, 1]$ such that $\sum_{i=1}^k p_i = 1$ and $\mu = \sum_{i=1}^k p_i \cdot \mu_i$
- $s \xrightarrow{\tau}_c \mu$ if there exist transitions $s \xrightarrow{\tau} \mu_i$ for $i = 1, \ldots, k$, and real numbers $p_1, \ldots, p_k \in]0, 1]$ and $p_0 \in [0, 1]$ such that $\sum_{i=0}^k p_i = 1$ and $\mu = p_0 \cdot \delta_s + \sum_{i=1}^k p_i \cdot \mu_i$. We call p_0 the "skip" probability.

So, the relation $\xrightarrow{\tau}_c$ is defined just as $\xrightarrow{a}_c$, except that in addition to τ-transitions it admits to "skip" performing any τ-transition with some probability.

With this special treatment of compound τ-transitions, a Dirac skip transition $s \xrightarrow{\tau}_c \delta_s$ is included for each state s. Furthermore, if $p_0 < 1$ then the compound transition $s \xrightarrow{\tau}_c p_0 \cdot \delta_s + \sum_{i=1}^k p_i \cdot \mu_i$ can be written as the convex combination of the Dirac skip transition $s \xrightarrow{\tau}_c \delta_s$ (with weight p_0) and the skip-free compound transition $s \xrightarrow{\tau} \sum_{i=1}^k p_i/(1-p_0)\mu_i$ (with weight $1-p_0$).

Weak internal transitions Weak transitions are a very easy-to-define concept in the non-probabilistic setting—weak τ-transitions simply arise as the reflexive and transitive closure of ordinary τ-transitions. In this respect, it is remarkable that the lifting of this concept to MDPs gets rather intricate. In the literature, there are (at least) three different approaches to define weak τ-transitions [8,10,20] (which in their effect are known to coincide *grosso modo* [4]). Intuitively, a weak τ-transition is obtained by scheduling sequences of ordinary τ-transitions in a way that almost surely terminates. We follow this idea using yet another notation, in which the relation $\Longrightarrow_c \subseteq S \times \mathrm{Dist}(S)$ comprising weak compound τ-transitions is generated from possibly infinite trees of compound τ-transitions.

Definition 2 *(Weak transitions and τ-trees)* A weak (internal compound) transition $s_0 \Longrightarrow_c \mu$ exists in a given MDP $(S, \mathrm{Act}, \rightarrow)$ iff there is a tree-like Markov chain $\mathcal{T} = (V, E, v_0, P_{\mathcal{T}}, \mathit{state})$—called a τ-tree in the sequel—where

- V denotes the set of nodes of $\mathcal{T}$,
- $E \subseteq V \times V$ is the edge relation,
- v_0 is the root of $\mathcal{T}$,
- $P_{\mathcal{T}} : E \rightarrow [0, 1]$ is the transition probability function, and
- $\mathit{state} : V \rightarrow S$ is a total function labelling tree nodes with MDP states, such that $\mathit{state}(v_0) = s_0$.

such that the following conditions (1) and (2) hold:

(1) the branching structure of each inner node v of $\mathcal{T}$ represents a compound τ-transition in the MDP $\mathcal{M}$ in the sense that there exist distributions $\nu_0, \nu_1, \ldots, \nu_k \in \text{Dist}(S)$ and real numbers $p_1, \ldots, p_k \in]0, 1]$ and $p_0 \in [0, 1[$ such that $\sum_{i=0}^{k} p_i = 1$, $\nu_0 = \delta_{state(v)}$, $state(v) \xrightarrow{\tau} \nu_i$ for $i = 1, \ldots, k$, and

(1.1) *Childs*(v) (defined for each inner node $v \in V$ as the set $\{w \in V : (v, w) \in E\}$) can be partitioned into pairwise disjoint sets $V_0, V_1, \ldots, V_k$ where $V_0 = \varnothing$ if $p_0 = 0$;

(1.2) for each $i \in \{0, 1, \ldots, k\}$:
- $\text{supp}(\nu_i) = \{state(w) : w \in V_i\}$,
- $state(w) \neq state(w')$ for all nodes $w, w' \in V_i$ with $w \neq w'$,
- $P_{\mathcal{T}}(v, w) = p_i \cdot \nu_i(state(w))$ for each node $w \in V_i$;

(2) $\mu(s)$ equals the probability to reach the set *Leaves*(s) from the root in $\mathcal{T}$, where *Leaves*(s) denotes the set of leaves v in $\mathcal{T}$ with $state(v) = s$.

Notably, in condition (1), $state(v) \xrightarrow{\tau}_c p_0 \cdot \nu_0 + p_1 \cdot \nu_1 + \cdots + p_k \cdot \nu_k$ is a compound τ-transition with skip probability p_0, corresponding to a randomized scheduling decision opting with probability p_i for transition $state(v) \xrightarrow{\tau} \nu_i$—except for p_0 which is the stopping probability in $state(v)$. As μ is a distribution on S, condition (2) implies that $\mathcal{T}$ almost surely reaches a leaf. So, $\mathcal{T}$ might have infinite paths, but they occur with probability 0.

For the special case where $\mathcal{T}$ consists just of its root, we get $s \Rightarrow_c \delta_s$ for all states $s \in S$.

We say $\mathcal{T}$ is *compressed* if in condition (1)

- the distributions $\nu_1, \ldots, \nu_k$ are pairwise distinct,
- $p_0 < 1$, and
- $p_0 > 0$ implies that the unique child w of v with $V_0 = \{w\}$ is a leaf of $\mathcal{T}$.

It is easy to see that each weak transition has a compressed τ-tree. Moreover, each compressed τ-tree corresponds to a partial scheduler that schedules only τ-transitions with positive probability and where scheduling *stop* with probability < 1 corresponds to the skip option in compressed τ-trees.

Example 1 (Weak transitions and the need for skipping) Let $\mathcal{M}$ be an MDP with states t_0, s_0, s_1, s_2, u and action set $\text{Act} = \{a, b, c, \tau\}$ such that:

$$
\begin{array}{lll}
t_0 \xrightarrow{c} \delta_u & t_1 \xrightarrow{c} \delta_u & \\
t_0 \xrightarrow{\tau} \delta_{s_0} & t_1 \xrightarrow{\tau} \theta & \text{where } \theta(s_1) = \theta(s_2) = 1/2 \\
s_i \xrightarrow{\alpha} \delta_u & s_i \xrightarrow{\tau} \delta_u & \text{for } i = 0, 1, 2 \\
u \xrightarrow{b} \delta_u. & &
\end{array}
$$

Figure 1 depicts the MDP $\mathcal{M}$. Now let $\mathcal{M}[s_i]$ denote the sub-MDP containing the states reachable from s_i. Then, $\mathcal{M}[s_0]$, $\mathcal{M}[s_1]$ and $\mathcal{M}[s_2]$ are isomorphic and hence, any reasonable notion of bisimilarity should identify s_0, s_1, s_2. But then also t_0 and t_1 should be probabilistically bisimilar. Indeed, we should expect that the two t-states form an equivalence class, the three s-states an equivalence class and the u-state a singleton equivalence class.

State t_1 has a weak transition $t_1 \Rightarrow_c \mu$ where $\mu(s_1) = 1/2$ and $\mu(u) = 1/2$. The corresponding τ-tree $\mathcal{T}_1$ has four nodes: the root t_1 and s_1, s_2, u with the obvious state labels and $P_{\mathcal{T}_1}(t_1, s_1) = P_{\mathcal{T}_1}(t_1, s_2) = 1/2$ and $P_{\mathcal{T}_1}(s_2, u) = 1$. Nodes s_1 and u are leaves. The branching structure in the inner nodes t_1 and s_2 corresponds to the τ-transitions $t_1 \xrightarrow{\tau} \theta$ and $s_2 \xrightarrow{\tau} \delta_u$.

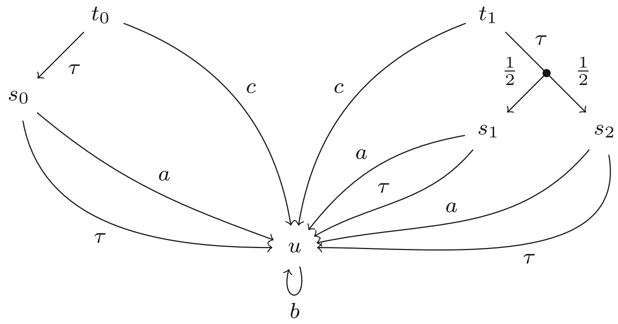

Fig. 1 Weak transition and the need for skipping

To match this weak transition $t_1 \Rightarrow_c \mu$, t_0 needs the skip option in state s_0. Indeed, we have $t_0 \Rightarrow_c \nu$ where $\nu(s_0) = \nu(u) = 1/2$. The corresponding τ-tree $\mathcal{T}_0$ has four nodes: v_0, v_1, v_2, v_3 where v_0 is the the root representing state t_0, its unique child v_1 with state label s_0, which again has two children v_2 and v_3 with state labels $state(v_2) = s_0$ and $state(v_3) = u$. Nodes v_2 and v_3 are leaves. The transition probabilities are $P_{\mathcal{T}_0}(v_0, v_1) = 1$ (obtained by the τ-transition from t_0) and $P_{\mathcal{T}_0}(v_1, v_2) = P_{\mathcal{T}_0}(v_1, v_3) = 1/2$ obtained by the compound τ-transition

$$s_0 \xrightarrow{\tau}_c \frac{1}{2}\delta_{s_0} + \frac{1}{2}\delta_u$$

from state $s_0 = state(v_1)$.

Compound and weak transitions over distributions The extensions of ordinary, compound and weak transitions to binary relations of distributions over states are defined as follows. Let $\mu, \nu \in \text{Dist}(S)$ and $\alpha \in \text{Act}$.

$$\mu \xrightarrow{\alpha}_c \nu$$

iff for each $s \in \text{supp}(\mu)$ there exists a compound α-transition $s \xrightarrow{\alpha}_c \nu_s$ such that:

$$\nu(t) = \sum_{s \in \text{supp}(\mu)} \mu(s) \cdot \nu_s(t) \quad \text{for all } t \in S$$

Weak transitions over distributions are defined as follows:

$$\mu \Rightarrow_c \nu$$

iff for each $s \in \text{supp}(\mu)$ there exists a weak transition $s \Rightarrow_c \nu_s$ such that:

$$\nu(t) = \sum_{s \in \text{supp}(\mu)} \mu(s) \cdot \nu_s(t) \quad \text{for all } t \in S$$

Obviously, $\Rightarrow_c$ as a binary relation over $\text{Dist}(S)$ is transitive, i.e., $\mu \Rightarrow_c \nu$ and $\nu \Rightarrow_c \theta$ implies $\mu \Rightarrow_c \theta$.

In the sequel, notations like $\mu \Rightarrow_c \xrightarrow{\alpha} \Rightarrow_c \nu$ are shorthand notations stating the existence of distributions μ', ν' such that $\mu \Rightarrow_c \mu' \xrightarrow{\alpha} \nu' \Rightarrow_c \nu$.

Countable MDPs Most results presented in this paper make no use of the default assumption that the underlying MDP is finite, they hold for countable MDPs, too. Countable MDPs are defined as in Definition 1, except that S, Act and $\rightarrow$ can be countable (possibly infinite) sets. Compound transitions in countable MDPs are defined in the same way, the only difference being that convex combinations of countably many transitions (and the skip option) are allowed. This propagates into the notions of τ-trees and weak transitions.

3 Operational definitions of bisimulations

The core of the bisimulation fragment of van Glabbeek's branching-time spectrum with silent moves [23] is formed by four behavioural relations that characterise in subtly different manners to what extent far the behaviours of two states in a labelled transition system are to be considered equal—in the presence of weak transitions. In what follows, we extend these definitions to the MDP setting [20,22].

Definition 3 *(Branching bisimulation)* An equivalence relation R is a *branching bisimulation* (or *b-bisimulation*) if for all $(s, t) \in R$, $\mu \in \mathrm{Dist}(S)$, and $\alpha \in \mathrm{Act}$:

(b) $s \xrightarrow{\alpha} \mu$ implies the existence of ν and ν' such that $t \Rightarrow_c \nu \xrightarrow{\alpha}_c \nu'$, $\delta_s \equiv_R \nu$, and $\mu \equiv_R \nu'$.

We write $s \approx_b t$ whenever there is a branching bisimulation R such that $(s, t) \in R$.

Notably, the reader may expect, as a parallel to the non-probabilistic case, in addition to condition (b) the alternative condition requiring that if $\alpha = \tau$, $s \xrightarrow{\tau} \mu$ implies that $\mu \equiv_R \delta_t$. However, due to the skip option, this condition is already considered as part of (b). Indeed, in particular if $s \xrightarrow{\tau} \mu$, it may be the case that $t \Rightarrow_c \delta_t \xrightarrow{\tau}_c \delta_t$ with $\delta_s \equiv_R \delta_t$, and $\mu \equiv_R \delta_t$, which is equivalent to the expected condition. A similar effect happens for the relations we define in the following.

Definition 4 (η*-bisimulation*) An equivalence relation R is an η*-bisimulation* if for all $(s, t) \in R$, $\mu \in \mathrm{Dist}(S)$, and $\alpha \in \mathrm{Act}$:

(η) $s \xrightarrow{\alpha} \mu$ implies the existence of ν and ν' such that $t \Rightarrow_c \nu \xrightarrow{\alpha}_c\Rightarrow_c \nu'$, $\delta_s \equiv_R \nu$, and $\mu \equiv_R \nu'$.

We write $s \approx_\eta t$ whenever there is an η-bisimulation R such that $(s, t) \in R$.

Definition 5 *(Delay bisimulation)* An equivalence relation R is a *delay bisimulation* (or *d-bisimulation*) if for all $(s, t) \in R$, $\mu \in \mathrm{Dist}(S)$, and $\alpha \in \mathrm{Act}$:

(d) $s \xrightarrow{\alpha} \mu$ implies the existence of $\nu \in \mathrm{Dist}(S)$ such that $t \Rightarrow_c\xrightarrow{\alpha}_c \nu$ and $\mu \equiv_R \nu$.

We write $s \approx_d t$ whenever there is a delay bisimulation R such that $(s, t) \in R$.

Definition 6 *(Weak bisimulation)* An equivalence relation R is a *weak bisimulation* (or *w-bisimulation*) if for all $(s, t) \in R$, $\mu \in \mathrm{Dist}(S)$, and $\alpha \in \mathrm{Act}$:

(w) $s \xrightarrow{\alpha} \mu$ implies that exists $\nu \in \mathrm{Dist}(S)$ such that $t \Rightarrow_c\xrightarrow{\alpha}_c\Rightarrow_c \nu$ and $\mu \equiv_R \nu$.

We write $s \approx_w t$ whenever there is a weak bisimulation R such that $(s, t) \in R$.

Divergence predicates The other important dimension in van Glabbeek's branching-time spectrum with silent moves [23] is spanned by different concepts in how far divergence is visible to an external observer. To resonate the distinctive aspects in the MDP setting, we define the state predicates Δ, s, 0, λ, and ∇^R, where R is an equivalence relation, as follows

$$
\begin{aligned}
\mathsf{s}(s) &\quad \text{iff} \quad s \overset{\tau}{\not\rightarrow}\,, \\
0(s) &\quad \text{iff} \quad s \overset{\alpha}{\not\rightarrow} \; \text{ for all } \alpha \in \text{Act (including } \tau), \\
\Delta(s) &\quad \text{iff} \quad \exists \sigma \text{ scheduler} : \Pr_s^\sigma\left((S \times \{\tau\})^\omega\right) = 1, \\
\lambda(s) &\quad \text{iff} \quad \Delta(s) \text{ or } 0(s), \\
\nabla^R(s) &\quad \text{iff} \quad \exists \sigma \text{ scheduler} : \Pr_s^\sigma\left(([s]_R \times \{\tau\})^\omega\right) = 1.
\end{aligned}
$$

Intuitively, $\mathsf{s}(s)$ means to say that "state s is stable"—it will not internally change behaviour — while $0(s)$ stands for "state s is dead"—it cannot exhibit any behavior. The predicate $\Delta(s)$ indicates that "state s can diverge" (almost surely), and $\lambda(s)$, quite obviously stands for "state s can diverge or is dead". Finally $\nabla^R(s)$ [16] indicates that "state s can diverge without altering behaviour"—assuming that equivalence relation R is equating behaviours in some way or another. In this case, we call state s an R-divergent state.

We lift these notions to distributions by defining for $\xi \in \{\nabla^R, \Delta, \mathsf{s}, 0, \lambda\}$ that $\xi(\mu)$ iff $\mu(\{s \mid \xi(s)\}) = 1$.

Behavioural equivalences respecting divergence Each of the above predicates can be combined with each of the relations we have seen before. The template is provided by the following definition.

Definition 7 We say that an equivalence relation R on S is ξ*-respecting*, for $\xi \in \{\Delta, \mathsf{s}, 0, \lambda\}$, if for all $(s, t) \in R$:

$$\xi(s) \text{ implies the existence of } \nu \in \text{Dist}(S) \text{ such that } t \Rightarrow_c \nu \text{ and } \xi(\nu).$$

We say that an equivalence relation R is ∇*-respecting* if for all $(s, t) \in R$:

$$\nabla^R(s) \text{ implies that } \nabla^R(t).$$

Note that compared to the strong preservation requirement of ∇-respecting equivalences, the requirement for ξ-respecting equivalences for $\xi \in \{\Delta, \mathsf{s}, 0, \lambda\}$ is more relaxed by permitting preservation of ξ up to weak transitions. For $\xi \in \{\mathsf{s}, 0, \lambda\}$, it is indeed possible that the ξ-predicate is not directly preserved by a ξ-respecting equivalence R (in the sense that $(s, t) \in R$, $\xi(s)$ and $\neg\xi(t)$ is possible). However, for the Δ-predicate, the strong preservation requirement is equivalent to the preservation of ξ up to weak transitions. Formally, if R is an equivalence on S then:

$$R \text{ is } \Delta\text{-respecting} \quad \text{iff} \quad (s, t) \in R \text{ and } \Delta(s) \text{ implies } \Delta(t)$$

The implication "$\Longleftarrow$" is obvious. To see why "$\Longrightarrow$" holds, consider a Δ-respecting equivalence R, two R-equivalent states s, t and suppose $\Delta(s)$. As R is Δ-respecting, state t has a weak transition $t \Rightarrow_c \nu$ such that $\Delta(u)$ for all states $u \in \text{supp}(\nu)$. Let $\mathcal{T}$ be a corresponding τ-tree. Pick schedulers σ_u for $u \in \text{supp}(\nu)$ such that $\Pr_u^{\sigma_u}((S \times \{\tau\})^\omega) = 1$. Let now σ be a scheduler which, when started in state t, first follows the decisions in the τ-tree $\mathcal{T}$ until having reached a state $u \in \text{supp}(\nu)$, upon which σ switches its mode and behaves as σ_u from that moment on. Obviously, we then have $\Pr_t^\sigma((S \times \{\tau\})^\omega) = 1$, and therefore $\Delta(t)$.

Induced ξ-respecting x-bisimilarity. Taking into account these last definitions, the relation $\approx_x^\xi$ for each $x \in \{b, \eta, d, w\}$ and $\xi \in \{\nabla, \Delta, \mathsf{s}, 0, \lambda, \varnothing\}$ is defined as the union of all ξ-respecting x-bisimulations. For example, $s \approx_\eta^{\mathsf{s}} t$ holds if there is an s-respecting η-bisimulation R such that $(s, t) \in R$. In particular $\approx_x^\varnothing$ means $\approx_x$.

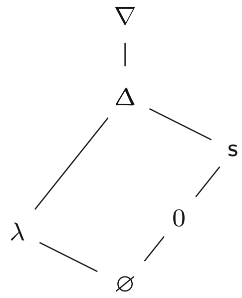

Fig. 2 Sensitivity of the divergence predicates

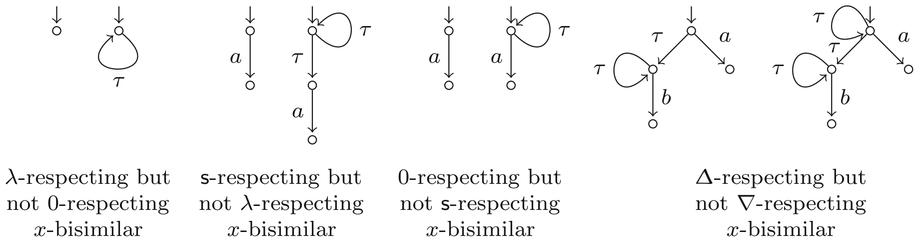

Fig. 3 Distinguishing examples for the different divergence notions

To simplify the notations for the induced equivalences on distributions, we will in the sequel write $\mu \approx_s^\xi \nu$ rather than $\mu \equiv_{\approx_s^\xi} \nu$.

As expected, each notion of divergence has a different distinguishing power. Figure 2 orders each notion accordingly. In it, when two notions are connected by a line, the notion placed above is more distinguishing than the one placed below. Such relation will be established in Sect. 8. Here, we only show that the distinguishing power is strict. Indeed, examples in Fig. 3 (the first three are borrowed from [23]) witness all possible differences. They should be understood jointly with Fig. 2. Thus, the first example in Fig. 3 provides two MDPs that are ξ_1-respecting x-bisimilar, but not ξ_2-respecting x-bisimilar, with $\xi_1 \in \{\lambda, \varnothing\}$ and $\xi_2 \in \{\nabla, \Delta, \mathsf{s}, 0\}$, and similarly for the other examples.

4 Basic properties of bisimulations

This section establishes some basic insights regarding the various predicates and bisimulation relations and their combinations. Several of the observations have appeared in one way or another at various places in the scientific literature, but we collect and prove them here for the sake of uniformity of presentation. All properties in this section apply to countably infinite MDPs.

We start with a trivial property regarding the liftings of equivalences on the state space to equivalences on distributions.

Lemma 1 *Let R be an equivalence relation on S and $\mu_1, \dots, \mu_k, \nu_1, \dots, \nu_k \in \mathrm{Dist}(S)$ such that $\mu_i \equiv_R \nu_i$ for $i = 1, \dots, k$, and let $p_1, \dots, p_k \in [0, 1]$ such that $\sum_{i=1}^k p_i = 1$. Then, $\sum_{i=1}^k p_i \cdot \mu_i \equiv_R \sum_{i=1}^k p_i \cdot \nu_i$.* □

Lemma 2 (Simulating weak transitions by bisimilar states) *Let R be a ξ-respecting x-bisimulation and $s, t \in S$, $\mu \in \mathrm{Dist}(S)$. Then:*

$$s \Rightarrow_c \mu \text{ and } (s,t) \in R \text{ implies } t \Rightarrow_c \nu \text{ and } \mu \equiv_R \nu \text{ for some } \nu \in \mathrm{Dist}(S).$$

Proof The first part of the proof spells out why it is no restriction to prove Lemma 2 for the case where all states in $\mathrm{supp}(\mu)$ and their R-equivalent states are terminal (see assumption (*) below). This assumption will ease the (rather technical) arguments provided in the second part.

Part 1. We provide a formal justification for the following assumption (*):

(*) All states in $\{u \in S \mid \exists u' \in \mathrm{supp}(\mu) \text{ such that } (u,u') \in R\}$ are terminal.

Let $\mathcal{N}$ denote the MDP resulting from $\mathcal{M}$ by adding

- pairwise distinct fresh copies $u_{\mathcal{N}}$ for each state $u \in S$, and
- transitions $u \xrightarrow{\tau} \delta_{u_{\mathcal{N}}}$ for all states $u \in S$.

Thus, the state space of $\mathcal{N}$ is $S_{\mathcal{N}} = S \cup \{u_{\mathcal{N}} : u \in S\}$ and state $u_{\mathcal{N}}$ can only be reached from state $s \in S$ via a path fragment inside $\mathcal{M}$ from s to state u, followed by the fresh τ-transition from u to $u_{\mathcal{N}}$. Moreover, the states $u_{\mathcal{N}}$ are terminal in the new MDP $\mathcal{N}$. We then have $s \Rightarrow_c \mu_{\mathcal{N}}$ in $\mathcal{N}$ where $\mu_{\mathcal{N}}(u_{\mathcal{N}}) = \mu(u)$ for all $u \in \mathrm{supp}(\mu)$ and $\mu_{\mathcal{N}}(t) = 0$ for all states $t \in S$ of the original MDP. The τ-tree for $s \Rightarrow_c \mu_{\mathcal{N}}$ in $\mathcal{N}$ is obtained by adding edges $(v, v_{\mathcal{N}})$ for all leaves v in $\mathcal{T}$ with $\mathit{state}(v_{\mathcal{N}}) = \mathit{state}(v)$ and $P_{\mathcal{T}}(v, v_{\mathcal{N}}) = 1$. We define $R_{\mathcal{N}} \subseteq S_{\mathcal{N}} \times S_{\mathcal{N}}$ as follows:

$$R_{\mathcal{N}} = R \cup \big\{(s_{\mathcal{N}}, t_{\mathcal{N}}) : (s,t) \in R\big\}$$

It is easy to see that $R_{\mathcal{N}}$ is a ξ-respecting x-bisimulation.[1]

Having established that t has a weak transition $t \Rightarrow_c \nu_{\mathcal{N}}$ in $\mathcal{N}$ such that $\mu_{\mathcal{N}} \equiv_{R_{\mathcal{N}}} \nu_{\mathcal{N}}$, the remaining argument is as follows. By definition of $R_{\mathcal{N}}$ and using the $\equiv_{R_{\mathcal{N}}}$-equivalence of $\mu_{\mathcal{N}}$ and $\nu_{\mathcal{N}}$ we obtain:

$$\mathrm{supp}(\nu_{\mathcal{N}}) \subseteq \big\{u_{\mathcal{N}} : u \in S\big\}$$

This yields $t \Rightarrow_c \nu$ in $\mathcal{M}$ where ν is the distribution on S given by $\nu(u) = \nu\big(u_{\mathcal{N}}\big)$ for all states $u \in S$. Moreover, $\mu \equiv_R \nu$ as $\mu_{\mathcal{N}} \equiv_{R_{\mathcal{N}}} \nu_{\mathcal{N}}$ and for each $C \in S/R$, the set $C_{\mathcal{N}} = \{u_{\mathcal{N}} : u \in C\}$ is an $R_{\mathcal{N}}$-equivalence class. This completes the considerations why the above assumption (*) is no restriction.

Part 2. We now suppose that (*) holds and prove the statement of Lemma 2. Let $\mathcal{T}$ be the tree associated with the weak transition $s \Rightarrow_c \mu$. We may assume that $\mathcal{T}$ is compressed. Assumption (*) implies that there is no inner node of $\mathcal{T}$ such that $\mathit{state}(v) \in \mathrm{supp}(\mu)$. (As the states in the R-equivalence classes of some state in $\mathrm{supp}(\mu)$ are terminal, this yields that all nodes in $\mathcal{T}$ that are labelled with a state that is R-equivalent to some state in $\mathrm{supp}(\mu)$ need to be leaves. In particular, none of these nodes are expanded in $\mathcal{T}$ with the skip option for compound τ-transitions.)

[1] When dealing with $\xi \in \{\mathsf{s}, 0, \lambda\}$ then $u_1 \approx_x^{\xi} u_2$ in $\mathcal{N}$ is possible, although u_1 and u_2 may not be $\approx_x^{\xi}$-bisimilar in $\mathcal{M}$. This, however, is irrelevant for our purposes as we deal here only with the relation $R_{\mathcal{N}}$ which enjoys the property that if $u_1, u_2 \in S$ then $(u_1, u_2) \in R$ iff $(u_1, u_2) \in R_{\mathcal{N}}$. Furthermore, with R also $R_{\mathcal{N}}$ is ξ-respecting as all fresh states are terminal. Therefore, if $\xi \in \{\mathsf{s}, 0, \lambda\}$ then the ξ-predicate in $\mathcal{N}$ is $\xi_{\mathcal{N}} = \{u_{\mathcal{N}} : u \in S\}$, while for $\xi \in \{\Delta, \nabla\}$ the ξ-predicates in $\mathcal{N}$ and $\mathcal{M}$ agree in the sense that $\Delta_{\mathcal{N}} = \Delta_{\mathcal{M}}$ and $\nabla_{\mathcal{N}}^{R_{\mathcal{N}}} = \nabla_{\mathcal{M}}^{R}$.

If v is a node in $\mathcal{T}$ then we write $depth(v)$ for the length (number of transitions) in the unique path from the root v_0 of $\mathcal{T}$ to v. Let $\mathcal{T}_n$ denote the n-th truncation of $\mathcal{T}$ that results from $\mathcal{T}$ by removing all nodes v with $depth(v) > n$. Thus, the depth of all nodes in $\mathcal{T}_n$ is at most n. If $\mu_n \in \mathrm{Dist}(S)$ is the distribution induced by (the leaves of) $\mathcal{T}_n$ then

$$s \Rightarrow_c \mu_n \quad \text{and} \quad \mu = \lim_{n\to\infty} \mu_n$$

where the limit of the distribution μ_n is taken pointwise. Assumption (*) yields that $\mu_0(u) \leqslant \mu_1(u) \leqslant \mu_2(u) \leqslant \ldots$ and $\mu(u) = \sup_{n\geqslant 0} \mu_n(u)$ for all states $u \in S$ that are R-equivalent to some state in $\mathrm{supp}(\mu)$.[2]

By induction on n we obtain the existence of weak transitions $t \Rightarrow_c \nu_n$ with $\mu_n \equiv_R \nu_n$, and we can assure that the associated tree $\mathcal{T}'_n$ is compressed and satisfies the following conditions.

- $\mathcal{T}'_n$ is a subtree of $\mathcal{T}'_{n+1}$. More precisely, $\mathcal{T}'_{n+1}$ extends $\mathcal{T}'_n$ by replacing some leaves in $\mathcal{T}'_n$ with trees for "proper" weak transitions. (Here, "proper" means that at least one τ-transition of $\mathcal{M}$ is involved. So, these leaves in $\mathcal{T}'_n$ are inner nodes of $\mathcal{T}'_{n+1}$.)
- Each leaf v of $\mathcal{T}'_n$ where $state(v)$ is R-equivalent to some state in $\mathrm{supp}(\mu)$ is a leaf in $\mathcal{T}'_{n+1}$ too.
- For each R-equivalence class $C \in S/R$ where $C \cap \mathrm{supp}(\mu)$ is nonempty, we have that $\mu_n(C) = \nu_n(C)$ is the probability to reach a leaf v in $\mathcal{T}'_n$ where $state(v) \in C$.

In this, we use the fact that the states represented by inner nodes of $\mathcal{T}$ do not belong to the R-equivalence classes of states contained in $\mathrm{supp}(\mu)$.

This yields that for each state $u \in S$ where u is R-equivalent to some state in $\mathrm{supp}(\mu)$ the sequence $(\nu_n(u))_{n\geqslant 0}$ is monotonically increasing. Let $\nu : S \to [0, 1]$ defined by:

- $\nu(u) = \sup_{n\geqslant 0} \nu_n(u)$ for all $u \in S$ with $(u, u') \in R$ for some $u' \in \mathrm{supp}(\mu)$,
- $\nu(u) = 0$ for all other states $u \in S$.

Let us first show that ν is a distribution that is $\equiv_R$-equivalent to μ. Obviously, we then have $\nu(u) = \lim_{n\to\infty} \nu_n(u)$ for the states $u \in S$ that are R-equivalent to some state in $\mathrm{supp}(\mu)$. Since $\mu_n \equiv_R \nu_n$ and μ is the limit of the μ_n's, we obtain $\mu_n(C) = \nu_n(C)$ and therefore $\mu(C) = \nu(C)$ for each R-equivalence class $C \in S/R$. But then $\sum_{C\in S/R} \nu(C) = \sum_{C\in S/R} \mu(C) = 1$, too. This yields $\nu \in \mathrm{Dist}(S)$ and $\mu \equiv_R \nu$.

The remaining task is to prove $t \Rightarrow_c \nu$. Let $\mathcal{T}'$ denote the tree that results as the limit of the trees $\mathcal{T}'_n$. If v is a leaf of $\mathcal{T}'$ then there is some $n \in \mathbb{N}$ such that v is a leaf in the trees $\mathcal{T}'_m$ for all $m \geqslant n$, and the probability to reach v in $\mathcal{T}'$ from its root equals the probability to reach v in $\mathcal{T}'_m$ from the root for each $m \geqslant n$. Vice versa, each leaf v of $\mathcal{T}'_n$ where $state(v)$ is R-equivalent to some state in $\mathrm{supp}(\mu)$ is a leaf of $\mathcal{T}'$. (This is because these states are terminal in $\mathcal{M}$ by assumption (*) and because $\mathcal{T}'$ is compressed.) Hence, for each state $u \in S$, $\nu(u)$ equals the probability to reach a leaf v with $state(v) = u$ in $\mathcal{T}'$ from its root. But then the probability to reach a leaf in $\mathcal{T}'$ is $\sum_{u\in S} \nu(u) = 1$. Hence, $\mathcal{T}'$ induces a weak transition $t \Rightarrow_c \nu$. ◻

Corollary 1 (Simulating weak transitions by bisimilar distributions) *Let R be as in Lemma 2 and let $\rho, \theta, \mu \in \mathrm{Dist}(S)$. Then:*

$$\rho \equiv_R \theta \text{ and } \rho \Rightarrow_c \mu \quad \text{implies} \quad \begin{cases} \text{there exists } \nu \in \mathrm{Dist}(S) \text{ with} \\ \theta \Rightarrow_c \nu \text{ and } \mu \equiv_R \nu \end{cases}$$

[2] Note, however, that $\mathrm{supp}(\mu_n)$ contains states $u \in S$ that do not have R-equivalent states in $\mathrm{supp}(\mu)$, unless $\mathcal{T}_n = \mathcal{T}$. These are exactly the states that occur as labels of some inner node v of $\mathcal{T}$ with $depth(v) = n$. For these states u, we have $\lim_{n\to\infty} \mu_n(u) = 0$, but no monotonicity property can be guaranteed for the sequence $(\mu_n(u))_{n\geqslant 0}$.

Proof Distribution μ has the form

$$\mu = \sum_{s \in \mathrm{supp}(\rho)} \rho(s) \cdot \mu_s$$

where $s \Rightarrow_c \mu_s$ for all states s in the support of ρ. Let $E \in S/R$ and $s \in E \cap \mathrm{supp}(\rho)$. By Lemma 2, each state $u \in E$ has a weak transition $u \Rightarrow_c \mu_{u,s}$ with $\mu_{u,s} \equiv_R \mu_s$. Let

$$\nu_u = \sum_{s \in E \cap \mathrm{supp}(\rho)} \frac{\rho(s)}{\rho(E)} \cdot \mu_{u,s}$$

For each $C \in S/R$ we have $\mu_{u,s}(C) = \mu_s(C)$. Hence, for each state $u \in E$:

$$\nu_u(C) = \sum_{s \in E \cap \mathrm{supp}(\rho)} \frac{\rho(s)}{\rho(E)} \cdot \mu_{u,s}(C) = \sum_{s \in E \cap \mathrm{supp}(\rho)} \frac{\rho(s)}{\rho(E)} \cdot \mu_s(C)$$

Therefore, $\nu_u(C) = \nu_{u'}(C)$ for all states $u, u' \in E$. Moreover:

$$\nu_u(C) \cdot \rho(E) = \sum_{s \in E \cap \mathrm{supp}(\rho)} \rho(s) \cdot \mu_s(C)$$

As $\rho \equiv_R \theta$ we have:

$$\sum_{u \in E} \rho(u) = \rho(E) = \theta(E) = \sum_{u \in E} \theta(u)$$

As the values $\nu_u(C)$ only depend on C, but not on u (see above), we obtain:

$$\sum_{u \in E} \theta(u) \cdot \nu_u(C) = \sum_{s \in E \cap \mathrm{supp}(\rho)} \rho(s) \cdot \mu_s(C)$$

Consider now the distribution ν defined as follows:

$$\nu = \sum_{u \in \mathrm{supp}(\theta)} \theta(u) \cdot \nu_u$$

We then have $\theta \Rightarrow_c \nu$. Let $C \in S/R$. Then:

$$\begin{aligned}
\mu(C) &= \sum_{s \in \mathrm{supp}(\rho)} \rho(s) \cdot \mu_s(C) \\
&= \sum_{E \in S/R} \sum_{s \in E \cap \mathrm{supp}(\rho)} \rho(s) \cdot \mu_s(C) \\
&= \sum_{E \in S/R} \sum_{u \in E} \theta(u) \cdot \nu_u(C) \\
&= \sum_{u \in \mathrm{supp}(\theta)} \theta(u) \cdot \nu_u(C) = \nu(C)
\end{aligned}$$

This yields $\mu \equiv_R \nu$. □

In Sect. 2, x-bisimulations have been introduced as binary relations of states. The constraints of an x-bisimulation R can now be lifted for the induced relation $\equiv_R$ on distributions as stated in the following lemma:

Lemma 3 (Bisimulation conditions for distributions) *Let R be an x-bisimulation where $x \in \{b, \eta, d, w\}$. Then, for all distributions $\rho, \theta, \mu \in \mathrm{Dist}(S)$ and all actions $\alpha \in \mathrm{Act}$ such that $\rho \xrightarrow{\alpha}_c \mu$ and $\rho \equiv_R \theta$:*

$$
\begin{array}{lll}
\textit{If } x = b & \textit{then} & \exists \theta', \nu \in \mathrm{Dist}(S) \textit{ s.t. } \theta \Rightarrow_c \theta' \xrightarrow{\alpha}_c \nu, \rho \equiv_R \theta' \textit{ and } \mu \equiv_R \nu. \\
\textit{If } x = \eta & \textit{then} & \exists \theta', \nu \in \mathrm{Dist}(S) \textit{ s.t. } \theta \Rightarrow_c \theta' \xrightarrow{\alpha}_c \Rightarrow_c \nu, \rho \equiv_R \theta' \textit{ and } \mu \equiv_R \nu. \\
\textit{If } x = d & \textit{then} & \exists \nu \in \mathrm{Dist}(S) \textit{ s.t. } \theta \Rightarrow_c \xrightarrow{\alpha}_c \nu \textit{ and } \mu \equiv_R \nu. \\
\textit{If } x = w & \textit{then} & \exists \nu \in \mathrm{Dist}(S) \textit{ s.t. } \theta \Rightarrow_c \xrightarrow{\alpha}_c \Rightarrow_c \nu \textit{ and } \mu \equiv_R \nu.
\end{array}
$$

Proof We provide the proof for $x = b$. The argument for $x \in \{\eta, d, w\}$ is analogous and omitted here.

We suppose $\rho \xrightarrow{\alpha}_c \mu$. Hence, for each state $s \in \mathrm{supp}(\rho)$ there exist compound transitions $s \xrightarrow{\alpha}_c \mu_s$ such that

$$\mu = \sum_{s \in \mathrm{supp}(\rho)} \rho(s) \cdot \mu_s$$

Being a compound transition of s indicates the existence of finite index sets I_s and (non-compound) transitions $s \xrightarrow{\alpha} \mu_{s,i}$ for $i \in I_s$ and real numbers $q_{s,i}$ in the interval $[0, 1]$ such that $\sum_{i \in I_s} q_{s,i} = 1$ and

$$\mu_s = \sum_{i \in I_s} q_{s,i} \cdot \mu_{s,i}$$

In case $\alpha = \tau$, the compound α-transition of s might use the skip option, i.e., there might be an index $i \in I_s$ such that $\mu_{s,i} = \delta_s$. Let us now fix a state $s \in \mathrm{supp}(\rho)$. For each triple (t, s, i) where $s \in \mathrm{supp}(\rho)$, $t \in [s]_R$ and $i \in I_s$ there exist distributions $\theta'_{t,s,i}$, $\nu_{t,s,i}$ such that

$$t \Rightarrow_c \theta'_{t,s,i} \xrightarrow{\alpha}_c \nu_{t,s,i} \quad \text{and} \quad \theta'_{t,s,i} \equiv_R \delta_s \quad \text{and} \quad \nu_{t,s,i} \equiv_R \mu_{s,i}$$

Note that if $\alpha = \tau$ and $\mu_{s,i} = \delta_s$ then we may deal with $\theta'_{t,s,i} = \delta_t$. Let now

$$\theta'_{t,s} = \sum_{i \in I_s} q_{s,i} \cdot \theta'_{t,s,i} \quad \text{and} \quad \nu_{t,s} = \sum_{i \in I_s} q_{s,i} \cdot \nu_{t,s,i}$$

Then, $t \Rightarrow_c \theta'_{t,s} \xrightarrow{\alpha}_c \nu_{t,s}$ and $\theta'_{t,s} \equiv_R \delta_s$ and $\nu_{t,s} \equiv_R \mu_s$.

Let now $E \in S/R$. If $\mathrm{supp}(\rho) \cap E$ is nonempty then for each $t \in E$ we define:

$$\theta'_t = \sum_{s \in \mathrm{supp}(\rho) \cap E} \frac{\rho(s)}{\rho(E)} \cdot \theta'_{t,s} \quad \text{and} \quad \nu_t = \sum_{s \in \mathrm{supp}(\rho) \cap E} \frac{\rho(s)}{\rho(E)} \cdot \nu_{t,s}$$

Furthermore, let

$$\mu_E = \sum_{s \in \mathrm{supp}(\rho) \cap E} \frac{\rho(s)}{\rho(E)} \cdot \mu_s$$

For each $t \in E$ we have: $t \Rightarrow_c \theta'_t \xrightarrow{\alpha}_c \nu_t$ and $\theta'_{t,s} \equiv_R \delta_s$ and $\nu_t \equiv_R \mu_E$.

Let $\mathfrak{E}$ denote the set of R-equivalence classes $E \in S/R$ such that $\mathrm{supp}(\rho) \cap E$ is nonempty. Then

$$\mu = \sum_{E \in \mathfrak{E}} \rho(E) \cdot \mu_E$$

We define:

$$\theta'_E = \sum_{t \in E} \frac{\theta(t)}{\theta(E)} \cdot \theta_t \quad \text{and} \quad \nu_E = \sum_{t \in E} \frac{\theta(t)}{\theta(E)} \cdot \nu_t$$

We then have $\theta'_E \equiv_R \delta_s$ and $\nu_E \equiv_R \mu_E$. Finally, we put

$$\theta' = \sum_{E \in \mathfrak{E}} \rho(E) \cdot \theta'_E \quad \text{and} \quad \nu' = \sum_{E \in \mathfrak{E}} \rho(E) \cdot \nu_E$$

We then have $\theta \Rightarrow_c \theta' \xrightarrow{\alpha}_c \nu$ and $\theta' \equiv_R \rho$ and $\mu \equiv_R \nu$. □

Stutter lemma and its consequences

The following considerations on "stutter steps" (i.e., invisible transitions within the same bisimulation equivalence class) will be useful for reasoning about branching and η-bisimulations as well as for ∇-divergence in the context of x-bisimulations for any $x \in \{b, \eta, w, d\}$.

In the non-probabilistic setting, it is well-known that all states that belong to a cycle built by τ-transitions are x-bisimilar (for any $x \in \{b, \eta, w, d\}$). Even stronger, whenever in a (non-probabilistic) labelled transition system R is an x-bisimulation and state s has a weak transition to some state t (i.e., a path from s to t built by τ-transitions) where t belongs to the R-equivalence class of s then all intermediate states in that weak transition from s to t are x-bisimilar to s, and so are the states that belong to the R-equivalence classes of the intermediate states. The following considerations serve to establish a corresponding result in the probabilistic setting by stating that, if R is an x-bisimulation, then all states that are R-equivalent to some state in the τ-tree $\mathcal{T}$ of some weak transition $s \Rightarrow_c \mu$ where $\mu \equiv_R \delta_s$ are x-bisimilar to s. (This will be a consequence of Lemma 6 stated below.)

To establish that result, we will first introduce the so-called *stutter extension* $\lhd_R$ of an x-bisimulation R, which extends R by the pairs (s, t) where t is R-equivalent to some state in the τ-tree for some weak transition $s \Rightarrow_c \mu$ where $\mu \equiv_R \delta_s$. The state-pairs (s, t) in $\lhd_R$ satisfy a cycle condition in the sense that there exists a scheduler that reaches the R-equivalence class of s from t via τ-transitions with probability 1, and vice versa. This will be shown in Lemma 4, which again will be used to show that $\lhd_R$ is itself an x-bisimulation, and that it is ξ-respecting if so is R (see Lemma 6 below, which will be referred to as the *stutter lemma*.) The stutter lemma will be a useful technical vehicle in the proofs of various statements regarding ∇-divergence. Among others, the stutter lemma permits the assumption that a given ξ-respecting x-bisimulation R is stutter-closed in the sense that all states that belong to the τ-tree of a weak transition $s \Rightarrow_c \mu$ where $\mu \equiv_R \delta_s$ are R-equivalent to s. Such weak transitions consist of transitions inside the same R-equivalence class, which corresponds to the classical notion of stuttering [12].

Stutter extensions Let R be an x-bisimulation for some $x \in \{b, \eta, d, w\}$. The stutter extension of R, denoted $\lhd_R$, is a binary relation on the state space S of the given MDP $\mathcal{M}$ defined as follows. Given states $s, t \in S$ then it holds:

$$s \lhd_R t$$

if and only if there exists a weak transition $s \Rightarrow_c \mu$ with associated τ-tree $\mathcal{T}$ such that the following conditions (i) and (ii) are satisfied:

(i) $\mu \equiv_R \delta_s$,

(ii) There is a node v in $\mathcal{T}$ such that $state(v) \in [t]_R$.

The following lemma shows that $s \lhd_R t$ implies that s has a weak transition to a distribution consisting of states in $[t]_R$, and vice versa.

Lemma 4 (Basic properties of stutter extensions) *With the notations as above, if $s \lhd_R t$ then:*

(a) *There exists $\mu_{s,t} \in \mathrm{Dist}(S)$ with $t \Rightarrow_c \mu_{s,t}$ and $\mu_{s,t} \equiv_R \delta_s$.*
(b) *There exists $\nu_{s,t} \in \mathrm{Dist}(S)$ with $s \Rightarrow_c \nu_{s,t}$ and $\nu_{s,t} \equiv_R \delta_t$.*

Proof Let $s \Rightarrow_c \mu$ where $\mu \equiv_R \delta_s$ and let $\mathcal{T}$ be a corresponding τ-tree and v a node of $\mathcal{T}$ with $state(v) \in [t]_R$.

For the proof of statement (a), we consider the sub-tree $\mathcal{T}[v]$ of $\mathcal{T}$ with root v. Then, $\mathcal{T}[v]$ is a τ-tree inducing a weak transition $state(v) \Rightarrow_c \rho_s$ where $\rho_s \equiv_R \delta_s$. As $state(v)$ and t are R-equivalent, Lemma 2 yields the existence of a distribution $\mu_{s,t}$ with $t \Rightarrow_c \mu_{s,t}$ and $\mu_{s,t} \equiv_R \delta_s$.

We now turn to the proof of statement (b). Let $T = [t]_R$. We may assume that $s \notin T$ (otherwise the claim is obvious). By considering the sub-tree $\mathcal{T}'$ of $\mathcal{T}$ resulting from $\mathcal{T}$ by turning all inner nodes v' with $state(v') \in T$ into a leaf (and removing their sub-trees) we obtain $s \Rightarrow_c \nu$ where $\mathrm{supp}(\nu) \subseteq \mathrm{supp}(\mu) \cup T$ and $\nu(T) > 0$ and such that no inner node of $\mathcal{T}'$ belongs to T. Let σ be a (randomized) scheduler inducing $s \Rightarrow_c \nu$ and an analogous weak transition $s' \Rightarrow_c \nu'$ with $\nu \equiv_R \nu'$ from all other states $s' \in [s]_R$. In particular, $p \stackrel{\text{def}}{=} \nu(T) = \nu'(T) > 0$ and $\nu([s]_R) = \nu'([s]_R) = 1-p$.

We now define a new scheduler σ^*. Starting in state s, σ^* behaves as σ until having generated a complete path in the tree associated with $s \Rightarrow_c \nu$, in which case the current state s' either belongs to T (which happens with probability p) or to $[s]_R$ (which happens with probability $1-p$). In the latter case, i.e., if $s' \in [s]_R$, scheduler σ^* repeats the procedure and mimicks σ from s' until having generated a complete path in the tree associated with $s' \Rightarrow_c \nu'$, in which case the current state s'' either belongs to T (which happens with probability p) or to $[s]_R$. And so on. Then, the probability under σ^* to reach T from s via τ-transitions is

$$\sum_{i=0}^{\infty} p(1-p)^i = 1.$$

This yields the existence of a weak transition $s \Rightarrow_c \nu_{s,t}$ where $\nu_{s,t}(T) = 1$, i.e., $\nu_{s,t} \equiv_R \delta_t$. □

We now use Lemma 4 to show the symmetry and transitivity of $\lhd_R$. (Recall that we use the notion “coarser” in the sense of “strictly coarser or equal”.)

Lemma 5 (Equivalence/refinement property of stutter extensions) *If R is an x-bisimulation then $\lhd_R$ is an equivalence relation and coarser than R.*

Proof The reflexivity of $\lhd_R$ is obvious.

The symmetry of $\lhd_R$ is a consequence of Lemma 4 and Corollary 1. Let us see why. Suppose $s \lhd_R t$. Let $\mu_{s,t}$ and $\nu_{s,t}$ be as in statements (a) and (b) of Lemma 4. Corollary 1 yields the existence of a distribution $\nu'_{s,t}$ such that $\mu_{s,t} \Rightarrow_c \nu'_{s,t}$ and $\nu'_{s,t} \equiv_R \nu_{s,t} \equiv_R \delta_t$.

Hence, $t \Rightarrow_c \mu_{s,t} \Rightarrow_c \nu'_{s,t}$. This yields a τ-tree for $t \Rightarrow_c \nu'_{s,t}$ that contains at least one node v where $state(v)$ that is R-equivalent to s. Hence, $t \lhd_R s$.

For the transitivity of $\lhd_R$ we suppose $s \lhd_R t$ and $t \lhd_R u$. Lemma 4 and Corollary 1 imply the existence of distributions ρ'_s, ρ'_t, ρ''_t, ρ'_u with

$$s \Rightarrow_c \rho'_t \Rightarrow_c \rho'_u \Rightarrow_c \rho''_t \Rightarrow_c \rho'_s$$

and $\rho_s' \equiv_R \delta_s$, $\rho_u' \equiv_R \delta_u$ and $\rho_t'' \equiv_R \rho_t' \equiv_R \delta_t$.

Using the notations of Lemma 4:

- $\rho_t' = \nu_{s,t}$.
- Distribution ρ_u' is then obtained by Corollary 1 applied to the $\equiv_R$-equivalent distributions δ_t and $\rho_t' = \nu_{s,t}$ and the weak transition $\delta_t \Rightarrow_c \nu_{t,u}$.
- Distribution ρ_t'' is obtained by Corollary 1 applied to the $\equiv_R$-equivalent distributions δ_u and $\nu_{t,u}$ and the weak transition $\delta_u \Rightarrow_c \mu_{t,u}$.
- Finally, distribution ρ_s' is obtained by Corollary 1 applied to the $\equiv_R$-equivalent distributions δ_t and ρ_t'' and the weak transition $\delta_t \Rightarrow_c \mu_{s,t}$.

Hence, $s \Rightarrow_c \rho_s'$ and there is a corresponding τ-tree that contains at least one node whose state belongs to $[u]_R$. This yields $s \lhd_R u$.

Finally, we prove that $\lhd_R$ is coarser than R. This is a direct consequence of the fact that a tree consisting of a root labeled with state s is a τ-tree for $s \Rightarrow_c \delta_s$. This yields $s \lhd_R t$ for all states $t \in [s]_R$ and therefore $R \subseteq \lhd_R$. ◻

We are now ready to prove that the stutter extension of ξ-respecting x-bisimulations are ξ-respecting x-bisimulations too.

Lemma 6 (Stutter lemma) *Let $x \in \{b, \eta, d, w\}$ and $\xi \in \{\nabla, \Delta, \mathsf{s}, 0, \lambda, \varnothing\}$. If R is a ξ-respecting x-bisimulation then so is $\lhd_R$. In particular, $s \lhd_R t$ implies $s \approx_x^\xi t$.*

Proof We first show the simulation condition. We provide the proof for $x = b$. The arguments for $x \in \{\eta, d, w\}$ are analogous and omitted here.

Suppose $s \lhd_R t$ and $s \xrightarrow{\alpha} \mu$. Part (a) of Lemma 4 yields the existence of a distribution ρ_s with $\rho_s \equiv_R \delta_s$ and $t \Rightarrow_c \rho_s$. Hence, $\mathrm{supp}(\rho_s) \subseteq [s]_R$. But then for each state $u \in \mathrm{supp}(\rho_s)$ there are distributions ν_u, ν_u' such that $u \Rightarrow_c \nu_u \xrightarrow{\alpha}_c \nu_u'$ and $\nu_u \equiv_R \delta_s$ and $\nu_u' \equiv_R \mu$. With

$$\nu = \sum_{u \in \mathrm{supp}(\rho_s)} \rho_s(u) \cdot \nu_u \quad \text{and} \quad \nu' = \sum_{u \in \mathrm{supp}(\rho_s)} \rho_s(u) \cdot \nu_u'$$

we get $t \Rightarrow_c \rho_s \Rightarrow_c \nu \xrightarrow{\alpha}_c \nu'$ and $\nu \equiv_R \delta_s$ and $\nu' \equiv_R \mu$.

Next we show that $\lhd_R$ is ξ-respecting. We start with the explicit divergence predicate $\xi = \nabla$. Suppose $s \lhd_R t$ where s is $\lhd_R$-diverging (i.e., $\nabla^{\lhd_R}(s)$ holds). The task is to prove that t is $\lhd_R$-diverging. Let C denote the $\lhd_R$-equivalence class of s (and t). Pick a scheduler σ such that all σ-paths from some state in $[s]_R$ consist of τ-steps inside C. (Note that such a scheduler exists as R is ∇-respecting and therefore $\nabla^R(s')$ for all states $s' \in C$.) Furthermore, we pick a τ-tree $\mathcal{T}$ for a weak transition $s \Rightarrow_c \mu$ where $\mu \equiv_R \delta_s$ such that $\mathcal{T}$ contains a node v with $\mathit{state}(v) \in [t]_R$. As in the proof of part (a) of Lemma 4, we regard the subtree $\mathcal{T}[v]$ of $\mathcal{T}$ with root v. It yields a weak distribution $\rho_s = \mu_{s,t}$ with $\rho_s \equiv_R \delta_s$. Consider now the following scheduler σ', which when started in state t, first follows the decisions in $\mathcal{T}[v]$. As soon as a leaf of $\mathcal{T}[v]$ has been reached then σ' behaves as σ. As the states of all nodes in $\mathcal{T}[v]$ belong to C, all σ'-paths from t are infinite τ-paths consisting of C-states. This yields that t is $\nabla^{\lhd_R}$-diverging.

Let us now consider $\xi \in \{\Delta, \mathsf{s}, 0, \lambda\}$ and suppose that $\xi(s)$ and $s \lhd_R t$. The task is to show that t has a weak transition $t \Rightarrow_c \nu$ such that $\xi(u)$ for all states $u \in \mathrm{supp}(\nu)$ (see Definition 7).

- We first discuss the cases $\xi \in \{\mathsf{s}, 0\}$. But then $\xi(s)$ implies that s has no outgoing τ-transition. Thus, $s \lhd_R t$ implies $t \in [s]_R$. As R is ξ-respecting, there is a weak transition $t \Rightarrow_c \nu$ such that $\xi(u)$ for all states $u \in \mathrm{supp}(\nu)$.

– Suppose now that $\xi = \Delta$. As $s \lhd_R t$, we can rely on part (a) of Lemma 4 to obtain a weak transition $t \Rightarrow_c \nu$ with $\nu \equiv_R \delta_s$. The latter yields $\mathrm{supp}(\nu) \subseteq [s]_R$. As R is Δ-respecting and $\Delta(s)$, we have $\Delta(u)$ for all states $u \in [s]_R$ (see the explanations after Definition 7). In particular, we have $\Delta(u)$ for all states $u \in \mathrm{supp}(\nu)$.
– As the λ-predicate is defined as the union of the Δ- and the 0-predicates, the case $\xi = \lambda$ follows with the same arguments as in the above two items.

□

Relation R is said to be *stutter-closed* if $R = \lhd_R$. A direct consequence of the stutter lemma is:

Corollary 2 (Existence of stutter-closed bisimulations) *Each ξ-respecting x-bisimulation is contained in some stutter-closed ξ-respecting x-bisimulation R.*

Proof This follows from the stutter lemma and the idempotence of the $\lhd$ construction if viewed as an operator on relations, i.e., $\lhd_{\lhd_R} = \lhd_R$. □

A consequence of Corollary 2 is that $\approx_x^\xi$ is the union of all stutter-closed ξ-respecting x-bisimulations.

Corollary 3 *Let $x \in \{b, \eta, w, d\}$ and let R be a stutter-closed x-bisimulation Then, for all states $s, t \in S$:*

$$(s,t) \in R \quad \textit{iff} \quad \begin{cases} \textit{there exist } \mu, \nu \in \mathrm{Dist}(S) s.t. \\ 1. \quad s \Rightarrow_c \nu \quad \textit{and} \quad \nu \equiv_R \delta_t \\ 2. \quad t \Rightarrow_c \mu \quad \textit{and} \quad \mu \equiv_R \delta_s \end{cases}$$

Proof The implication "$\Longrightarrow$" holds for any equivalence relation R as we may deal with $\nu = \delta_s$ and $\mu = \delta_t$.

For the implication "$\Longleftarrow$", let ν and μ be as in 1. and 2. Then, $\nu \equiv_R \delta_t$ implies that all states in $\mathrm{supp}(\nu)$ are R-equivalent to t. Hence, each state $u \in \mathrm{supp}(\nu)$ has a weak transition $u \Rightarrow_c \mu_u$ such that $\mu \equiv_R \mu_u$ (Lemma 2). As $\mu \equiv_R \delta_s$, this yields that the distributions μ_u are $\equiv_R$-equivalent to δ_s. This induces the existence of a distribution μ' with

$$s \Rightarrow_c \nu \Rightarrow_c \mu' \quad \text{and} \quad \mu' \equiv_R \delta_s$$

As $\nu \equiv_R \delta_t$ this implies the existence of a τ-tree for the weak transition $s \Rightarrow_c \mu'$ that contains a node labeled by a state that is R-equivalent to t. Hence, $s \lhd_R t$. As R is stutter-closed we have $R = \lhd_R$. This implies $(s,t) \in R$. □

Stutter transitions and stutter steps Let R be an equivalence relation on S. Transitions $s \xrightarrow{\tau} \mu$ where $\mu \equiv_R \delta_s$ are called R-stutter transitions. Note that $\mu \equiv_R \delta_s$ implies that $\mathrm{supp}(\mu)$ is contained in the R-equivalence class of state s. In what follows, we shall write $s \Longrightarrow_c^R \mu$ if $s \Rightarrow_c \mu$ in the sub-MDP $\mathcal{M}_R$ of $\mathcal{M}$ consisting of R-stutter transitions. Formally, if $\mathcal{M} = (S, \mathrm{Act}, \longrightarrow)$ then $\mathcal{M}_R = (S, \{\tau\}, \longrightarrow_R)$ where $s \xrightarrow{\tau}_R \mu$ iff $s \xrightarrow{\tau} \mu$ in $\mathcal{M}$ and $\mathrm{supp}(\mu) \subseteq [s]_R$. We refer to $s \Longrightarrow_c^R \mu$ as a R-stutter step.

As a consequence of the stutter lemma (Lemma 6) we obtain:

Corollary 4 *Let $x \in \{b, \eta, d, w\}$, $\xi \in \{\nabla, \Delta, s, 0, \lambda, \varnothing\}$ and let R be a stutter-closed ξ-respecting x-bisimulation, s a state and μ a distribution on states such that $s \Rightarrow_c \mu$ in the MDP $\mathcal{M}$ and $\delta_s \equiv_R \mu$. Then, $s \Longrightarrow_c^R \mu$.*

Corollary 4 yields alternative characterisations of branching and η-bisimulation. Namely:

Corollary 5 (Alternative characterisation of b-and η-bisimulations) *Let R be an equivalence relation on S. Then:*

- *R is a b-bisimulation iff for each $(s,t) \in R$ and each transition $s \xrightarrow{\alpha} \mu$ there exist $\nu, \nu' \in \mathrm{Dist}(S)$ such that $t \Longrightarrow_c^R \nu \xrightarrow{\alpha}_c \nu'$ with $\nu' \equiv_R \mu$.*
- *R is an η-bisimulation iff for each $(s,t) \in R$ and each transition $s \xrightarrow{\alpha} \mu$ there exist $\nu, \nu' \in \mathrm{Dist}(S)$ such that $t \Longrightarrow_c^R \nu \xrightarrow{\alpha}_c \Longrightarrow_c \nu'$ with $\nu' \equiv_R \mu$.*

Coarsest bisimulation property

Following the classical coinduction principle, we introduced the relation $\approx_x^\xi$ as the union of all ξ-respecting x-bisimulations. As such, $\approx_x^\xi$ is obvioulsly coarser than each of the ξ-respecting x-bisimulations. It is, however, not immediate that $\approx_x^\xi$ is again an equivalence relation and that it itself satisfies the conditions of being a ξ-respecting x-bisimulation. In particular, the transitivity of $\approx_x^\xi$ is not obvious in the probabilistic setting. For the sake of completeness, we present below a full proof for branching bisimilarity with different divergence predicates ξ (the relations $\approx_b^\xi$) and stress that analogous proof techniques are applicable for the relations $\approx_x^\xi$ where $x \in \{w, \eta, d\}$. Indeed, proofs of analogous statements for different (bi)simulation relations in MDP or other probabilistic models appeared in the literature, including Segala's work [20] who—among others—presented a detailed proof for the transitivity of forward simulations in MDPs or the seminal work on bisimulations for Markov processes on Polish spaces by Blute et al. [3] that uses advanced techniques of continuous mathematics to establish the transitivity of bisimilarity.

Lemma 7 (Coarsest bisimulation) *Let $x \in \{b, \eta, d, w\}$ and $\xi \in \{\nabla, \Delta, s, 0, \lambda, \varnothing\}$. Then, $\approx_x^\xi$ is the coarsest ξ-respecting x-bisimulation.*

Proof Clearly, it suffices to prove that $\approx_x^\xi$ is a ξ-respecting x-bisimulation.

Let $\approx$ denote the coarsest equivalence containing $\approx_x^\xi$. Obviously, $\approx_x^\xi$ is symmetric and reflexive (as all x-bisimulations are). Thus, $\approx$ is the transitive closure of $\approx_x^\xi$. We now show that $\approx$ is a ξ-respecting x-bisimulation which then yields that $\approx$ is contained in $\approx_x^\xi$. As $\approx$ subsumes $\approx_x^\xi$ (by definition), we then can conclude that $\approx$ equals $\approx_x^\xi$ and that $\approx_x^\xi$ is a ξ-respecting x-bisimulation.

We start with a simple general observation:

Claim 1. Let R and R' be equivalences such that R is coarser than R'. Then, $\equiv_R$ is coarser than $\equiv_{R'}$.

Proof of claim 1. As R is coarser than R', each R-equivalence class C can be written as a disjoint union of R'-equivalence classes, say $C = \bigcup_{i \in I} C_i$ where $C_i \in S/R'$ for each index $i \in I$. (The index set I is finite in finite MDPs and countable in the general case of countable MDPs.) But then $\theta(C) = \sum_{i \in I} \theta(C_i)$ for each $\theta \in \mathrm{Dist}(S)$. Hence, if θ_1 and θ_2 are distributions with $\theta_1 \equiv_{R'} \theta_2$ then $\theta_1(C) = \theta_2(C)$ for each $C \in S/R$, and therefore $\theta_1 \equiv_R \theta_2$.

Claim 2. $\approx$ is an x-bisimulation.

Proof of claim 2. We consider here only the case $x = b$. The proof of other cases is analogous and omitted here. So, let s, t be states in the given MDP $\mathcal{M}$ such that $s \approx t$ and suppose we are given a transition $s \xrightarrow{a} \mu$. As $\approx$ is the transitive closure of $\approx_x^\xi$ there exist a sequence of states $s_0, s_1, \ldots, s_n$ and a sequence $R_1, \ldots, R_n$ of ξ-repsecting b-bisimulations such that $s_0 = s$, $s_n = t$ and $(s_{k-1}, s_k) \in R_k$ for each $k \in \{1, \ldots, n\}$.

We now show by induction on $k \in \{0, 1, \ldots, n\}$ that there are distributions ν_k, ν'_k with

$$s_k \Rightarrow_c \nu_k \xrightarrow{\alpha}_c \nu'_k \text{ such that } \nu_k \equiv_\approx \delta_s \text{ and } \nu'_k \equiv_\approx \mu.$$

The basis of induction (k=0) is trivial as we may deal with $\nu_0 = \delta_s$ and $\nu'_0 = \mu$. In the induction step ($k-1 \Longrightarrow k$ where $1 \leqslant k \leqslant n$) we suppose that we are given distributions ν_{k-1} and ν'_{k-1} with $s_{k-1} \Rightarrow_c \nu_{k-1}$ where $\nu_{k-1} \equiv_\approx \delta_s$ and $\nu'_k \equiv_\approx \mu$. As $(s_{k-1}, s_k) \in R_k$ we can apply Lemma 2 to obtain the existence of a weak transition $s_k \Rightarrow_c \theta$ where $\nu_{k-1} \equiv_{R_k} \theta$. Applying part (a) of Lemma 3 to the b-bisimulation R_k and the compound transition $\nu_{k-1} \xrightarrow{\alpha}_c \nu'_{k-1}$ yields the existence of distributions ν_k, ν'_k such that

$$\theta \Rightarrow_c \nu_k \xrightarrow{\alpha}_c \nu'_k,\ \nu_k \equiv_{R_k} \nu_{k-1} \text{ and } \nu'_k \equiv_{R_k} \mu.$$

Hence, $s_k \Rightarrow_c \nu_k \xrightarrow{\alpha}_c \nu'_k$ and, using Claim 1, $\nu_k \equiv_\approx \nu_{k-1}$ and $\nu'_k \equiv_\approx \mu$.

With $\nu = \nu_n$, $\nu' = \nu'_n$ we obtain $t \Rightarrow_c \nu \xrightarrow{\alpha}_c \nu'$ where ν and ν' are distributions that are $\equiv_\approx$-equivalent to δ_s and μ, respectively.

Claim 3. $\approx$ is ξ-respecting.

Proof of claim 3. Again, we concentrate on the case $x = b$. Suppose s and t are states with $s \approx t$ and $\xi(s)$. There exists states $s_0, s_1, \ldots, s_n$ and stutter-closed ξ-respecting x-bisimulations $R_1, \ldots, R_n$ such that $s_0 = s$, $s_n = t$ and $(s_{k-1}, s_k) \in R_k$ for $k = 1, \ldots, n$. (For this, we use Corollary 2.)

Let us first consider $\xi \in \{\mathsf{s}, \Delta, 0, \lambda\}$ and let X denote the set of states u with $\xi(u)$. By assumption $s \in X$. The task is to show that there is a distribution μ with $t \Rightarrow_c \mu$ and $\mathrm{supp}(\mu) \subseteq X$. For this, we show by induction on k that there is a weak transition $s_k \Rightarrow_c \mu_k$ such that $\mathrm{supp}(\mu_k) \subseteq X$. The basis of induction ($k = 0$) is trivial as we can deal with $\mu_0 = \delta_s$. In the step of induction we suppose that $s_{k-1} \Rightarrow_c \mu_{k-1}$ and $\mathrm{supp}(\mu_{k-1}) \subseteq X$. By Lemma 2, state s_k has a weak transition $s_k \Rightarrow_c \theta$ with $\theta \equiv_{R_k} \nu_{k-1}$. In particular, all states in $\mathrm{supp}(\theta)$ are R_k-equivalent to some state in X. As R_k is ξ-respecting, each state $u \in \mathrm{supp}(\theta)$ has a weak transition $u \Rightarrow_c \rho_u$ such that $\mathrm{supp}(\rho_u) \subseteq X$. Let

$$\nu_k = \sum_{u \in \mathrm{supp}(\theta)} \theta(u) \cdot \rho_u$$

Then, $\theta \Rightarrow_c \nu_k$ and $\mathrm{supp}(\nu_k) = \bigcup_{u \in \mathrm{supp}(\rho_u)} \mathrm{supp}(\rho_u) \subseteq X$.

It remains to consider $\xi = \nabla$ and to show that $\approx$ is ∇-respecting. For this, we pick a $\approx$-diverging state s, i.e., a state s with $\nabla^\approx(s)$, and show that each state $t \in [s]_\approx$ is $\approx$-diverging. Pick states $s_0, s_1, \ldots, s_n$ and stutter-closed ∇-respecting x-bisimulations $R_1, \ldots, R_n$ such that $s_0 = s$, $s_n = t$ and $(s_{k-1}, s_k) \in R_k$ for $k = 1, \ldots, n$. (Note that Corollary 2 permits to suppose the stutter-closedness of $R_1, \ldots, R_n$.) By induction on $k \in \{0, \ldots, n\}$ we now show that s_k is $\approx$-diverging. The basis of induction is obvious as $s = s_0$ is $\approx$-diverging by assumption. In the step of induction ($k-1 \Longrightarrow k$ where $1 \leqslant k \leqslant n$) we suppose that s_{k-1} is $\approx$-diverging and aim to show that s_k is $\approx$-diverging too. If s_{k-1} is R_k-diverging then so is s_k and we are done. Suppose now that s_{k-1} is $\approx$-diverging, but not R_k-diverging. Then, the scheduler witnessing $\nabla^\approx(s_{k-1})$ induces a sequence of distributions $\mu_0, \mu_1, \mu_2, \ldots$ such that the following conditions hold:

- $\mu_0 = \delta_{s_{k-1}}$

- $\mu_0 \Longrightarrow_c^{R_k} \xrightarrow{\tau}_c \mu_1 \Longrightarrow_c^{R_k} \xrightarrow{\tau}_c \mu_2 \Longrightarrow_c^{R_k} \xrightarrow{\tau}_c \ldots$
- for each $i \in \mathbb{N}$:
  - if $\mu_i \equiv_{R_k} \mu_{i+1}$ then $\mu_i \equiv_{R_k} \mu_j$ for all $j > i$
  - $\mathrm{supp}(\mu_i) \subseteq [s_{k-1}]_\approx$

As s_{k-1} and s_k are R_k-equivalent, there is a sequence of distributions $\nu_0, \nu_1, \nu_2, \ldots$ such that:

- $\nu_0 = \delta_{s_k}$
- $\nu_0 \Longrightarrow_c^{R_k} \xrightarrow{\tau}_c \nu_1 \Longrightarrow_c^{R_k} \xrightarrow{\tau}_c \nu_2 \Longrightarrow_c^{R_k} \xrightarrow{\tau}_c \ldots$
- for each $i \in \mathbb{N}$: $\mu_i \equiv_{R_k} \nu_i$

As R_k is finer than $\approx$, we get $\mathrm{supp}(\nu_i) \subseteq [s_{k-1}]_\approx$ for all $i \in \mathbb{N}$. Hence, the sequence $(\nu_i)_{i\in\mathbb{N}}$ induces a scheduler witnessing $\nabla^\approx(s_k)$. □

For the coarsest ξ-respecting x-bisimulation $R = \approx_x^\xi$ we shall also write $s \Longrightarrow_c^{x,\xi} \mu$ rather than $s \Longrightarrow_c^R \mu$.

5 Properties of the divergence operators ∇ and Δ

We now turn our attention specifically to the ∇- and Δ-respecting bisimulations and to divergence probabilities. The latter refers to the set of probability values with which divergence may occur for some scheduler. As far as we are aware, the results we are going to establish below have not appeared elsewhere before. In contrast to Sect. 4 the results on divergence probabilities generally rely on the assumption that the MDP at hand is finite-state. We shall discuss in detail below why this is the case.

5.1 Explicit divergence

Recall that state s is said to be R-divergent if $\nabla^R(s)$ holds. We start with the following characterization of R-divergent states, which holds in any (possibly countably infinite) MDP:

Lemma 8 (Characterization of R-divergence) *Let $x \in \{b, \eta, d, w\}$ and let R be a stutter-closed (possibly not ∇-respecting) x-bisimulation and s a state. Then, the following statements are equivalent:*

(a) *s is R-divergent.*
(b) *There exists a subset W of $[s]_R$ such that $s \in W$ and each state $t \in W$ has a transition $t \xrightarrow{\tau} \mu_t$ with $\mathrm{supp}(\mu_t) \subseteq W$. (In particular, $t \xrightarrow{\tau} \mu_t$ is a R-stutter transition with $\mu_t \equiv_R \delta_s$.)*
(c) *There exists a scheduler σ such that*

$$\mathrm{Pr}_s^\sigma\left((T \times \{\tau\})^\omega\right) = 1$$

where T denotes the set of states t such that t has a weak transition $t \Rightarrow_c \mu$ with $\mu \equiv_R \delta_s$.

Proof Obviously, statement (a) implies (b). For this, we pick a scheduler σ witnessing the R-divergence of s and define W to be the set of states that occur on a σ-path from s.

Statement (c) is a consequence of statement (b). For this, we can regard any memoryless scheduler σ that schedules the τ-transition $t \xrightarrow{\tau} \mu_t$ for all states $t \in W$. As $s \in W \subseteq [s]_R \subseteq T$, all infinite σ-paths from s belong to $(T \times \{\tau\})^\omega$. Hence, the condition stated in (c) holds.

It remains to show that statement (c) implies (a). Let σ be a scheduler as in (c). By the definition of T, for each state $t \in T$ there is a distribution ν_t such that $t \Longrightarrow_c \nu_t$ and $\nu_t \equiv_R \delta_s$. Let T_n be the set of states that are reached from s via a σ-path of length n and let $\theta_n \in \mathrm{Dist}(T_n)$ be the distribution induced by σ. That is, $\theta_n(t)$ equals the probability to reach t from s in n steps under scheduler σ. Let $T_* = \bigcup_n T_n$. Then, T_* consists of all states that appear in a σ-path from s. Moreover, $T_* \subseteq T$. Then, $s \Longrightarrow_c \theta_n$ and $\theta_n \Longrightarrow_c \rho_n$ where

$$\rho_n = \sum_{t \in T_n} \theta_n(t) \cdot \nu_t$$

As $\nu_t \equiv_R \delta_s$ we have $\mathrm{supp}(\nu_t) \subseteq [s]_R$ for all $t \in T_n$, and therefore:

$$\mathrm{supp}(\rho_n) = \bigcup_{t \in T_n} \mathrm{supp}(\nu_t) \subseteq [s]_R$$

But then $\rho_n \equiv_R \delta_s$. This yields $s \lhd_R t$ for all states $t \in T_*$. By assumption, R is stutter-closed. That is, $R = \lhd_R$. But then $T_* \subseteq [s]_R$. Thus, σ is a witness for the R-divergence of s. □

As a consequence of the equivalence of statements (a) and (b) in Lemma 8 we get (again, no matter whether the underlying MDP is finite or infinite):

Corollary 6 (MD-schedulers witnessing explicit divergence) *Let $x \in \{b, \eta, d, w\}$ and R an x-bisimulation. Then, there is a memoryless deterministic scheduler σ_R such that for each state s:*

$$s \textit{ is } R-\textit{divergent iff } \mathrm{Pr}_s^{\sigma_R}\big([s]_R \times \{\tau\})^\omega\big) = 1$$

Proof For each R-divergent state t, let W_s be a subset of $[s]_R$ satisfying the condition of statement (b) in Lemma 8, and let $s \xrightarrow{\tau} \mu_s$ be a transition such that $\mathrm{supp}(\mu_s) \subseteq W_s$. Let now σ_R be a memoryless deterministic scheduler that assigns the transition $s \xrightarrow{\tau} \mu_s$ to each R-divergent s, an arbitrary transition $s \xrightarrow{\alpha} \mu_s$ to each non-terminal state s that is not R-divergent and *stop* to each terminal state. Then, all σ_R-paths starting in an R-divergent state are infinite and consist of τ-transitions inside $[s]_R$. This yields the claim. □

Explicit divergence probabilities We defined the predicate ∇^R as the set of R-divergent states by the existence of a scheduler σ where almost all σ-paths π from s are R-divergent in the sense that π consists of states in $[s]_R$ and τ-transitions. We now show that ∇-respecting bisimulations preserve precise divergence probabilities (see Lemma 11 below).

Let R be an x-bisimulation and s a state. We define $\mathbb{DP}_\nabla(s, R)$ as the set of R-divergence probability values of state s:

$$\mathbb{DP}_\nabla(s, R) = \Big\{\mathrm{Pr}_s^\sigma\big(([s]_R \times \{\tau\})^\omega\big) : \sigma \text{ is a scheduler for } \mathcal{M}\Big\}$$

Thus, $\mathbb{DP}_\nabla(s, R)$ is a nonempty subset of the real interval [0, 1]. Using standard results on extremal probability values for ω-regular properties in finite-state MDPs (see e.g. [1]), there are memoryless deterministic schedulers achieving the maximal resp. minimal probability for $([s]_R \times \{\tau\})^\omega$ and every state s. That is, in finite MDPs we have:

Lemma 9 (MD-schedulers with extremal ∇-divergence probabilities) *Let $x \in \{b, \eta, w, d\}$ and let R be an x-bisimulation in a finite MDP. Then, there exist memoryless deterministic schedulers $\sigma_R^{\max}$ and $\sigma_R^{\min}$ such that*

$$\mathrm{Pr}_u^{\sigma_R^{\max}}\big((S \times \{\tau\})^\omega\big) = \sup \mathbb{DP}_\nabla(u, R)$$

$$\Pr_u^{\sigma_R^{\min}}\big((S \times \{\tau\})^\omega\big) = \inf \mathbb{DP}_\nabla(u, R)$$

for all states $u \in S$. In particular,

$$\sup \mathbb{DP}_\nabla(s, R) = \max \mathbb{DP}_\nabla(s, R) \textit{ and } \inf \mathbb{DP}_\nabla(s, R) = \min \mathbb{DP}_\nabla(s, R).$$

Lemma 10 (Divergence probabilities for ∇-respecting bisimulations) *Let $x \in \{b, \eta, w, d\}$ and let R be a ∇-respecting x-bisimulation in a finite MDP and s a state. Then, the following statements are equivalent:*

(a) *s is R-divergent (i.e., $\nabla^R(s)$).*
(b) $\sup \mathbb{DP}_\nabla(s, R) = 1$
(c) $\sup \mathbb{DP}_\nabla(s, R) > 0$
(d) *There is a scheduler σ such that all maximal σ-paths from s belong to* $([s]_R \times \{\tau\})^\omega$.

Proof The equivalence of (a) and (b) is obvious where for (b) $\Longrightarrow$ (a) we use $\sup \mathbb{DP}_\nabla(s, R) = \max \mathbb{DP}_\nabla(s, R)$ as stated above. The implications (b) $\Longrightarrow$ (c) and (d) $\Longrightarrow$ (a) are trivial.

We prove the implication (a) $\Longrightarrow$ (d) by contraposition. So, we suppose that (d) does not hold and show that (a) does not hold. Hence, for each scheduler σ there is a maximal σ-path $\pi = s_0\, \alpha_1\, s_1\, \alpha_2\, s_2\, \alpha_3 \ldots$ with $s_0 = s$ and $\pi \notin ([s]_R \times \{\tau\})^\omega$. Let $n \in \mathbb{N}$ be the smallest index such that $s_n \notin [s]_R$ or $\alpha_n \neq \tau$, and let q be the probability under σ for generating the prefix $\pi' = s_0\, \alpha_1\, \ldots\, \alpha_n\, s_n$ of π. That is, $q = \Pr_s^\sigma(Cyl(\pi'))$ where $Cyl(\pi')$ denotes the cylinder set spanned by π'. The latter set consists of all maximal paths where π' is a prefix. Then, $q > 0$ and $Cyl(\pi') \cap ([s]_R \times \{\tau\})^\omega$ is empty. But then the divergence probability under σ from s is at most $1-q$. So, state s is not R-divergent, i.e., (a) does not hold.

We now turn to the proof of (c) $\Longrightarrow$ (a). As $\mathcal{M}$ is finite, we can rely on de Alfaro's observation [6], that under each scheduler σ, the limit of almost all σ-paths constitutes an end component of $\mathcal{M}$. Here, the limit of an infinite path π denotes the set of all states and transitions that appear infinitely often in π. An end component is a strongly connected sub-MDP.

Let $C = [s]_R$ and let $\mathfrak{E}$ denote the set of end components $\mathcal{E}$ of $\mathcal{M}$ that consist of states in C and τ-transitions. Then, $\Pr_s^\sigma\big((C \times \{\tau\})^\omega\big)$ equals the probability under σ to reach an end component $\mathcal{E} \in \mathfrak{E}$ from s along a path consisting of C-states and τ-transitions. Now, if $\sup \mathbb{DP}_\nabla(s, R)$ is positive (assumption (c)) then $\mathfrak{E}$ is nonempty. Pick some end component $\mathcal{E} \in \mathfrak{E}$. Obviously, all states in $\mathcal{E}$ are R-diverging. As $\mathcal{E}$ consists of C-states and R is ∇-respecting, state s is R-divergent too. □

Remark 1 (Divergence probabilities in infinite MDPs) By using arguments with end components, the proof of Lemma 10 heavily relies on the default assumption that the given MDP $\mathcal{M}$ is finite. Indeed, the characterization of R-divergence in Lemma 10 is wrong for countably infinite MDPs, even for the subclass of countable and finitely branching MDPs.

To illustrate this phenomenon, we construct on infinite MDP using a construction that shares similarities with an example provided [7] of a probabilistic process where probability mass of visible behavior is lost in divergence. Pick a sequence $(q_n)_{n \geqslant 0}$ of rational numbers q_n in the open interval $]0, 1[$ such that $\prod_{n=0}^{\infty} q_n$ converges to some positive value q. Thus, $0 < q < 1$. (For example, we can deal with $q_n = 1-1/2^n$.) Let now $\mathcal{M}$ be the MDP with state space $S = \{s_n : n \in \mathbb{N}\} \cup \{t, u\}$, action set $\mathrm{Act} = \{a, \tau\}$ and transitions

$$s_n \xrightarrow{\tau} \mu_n \text{ where } \mu_n(s_{n+1}) = q_n \text{ and } \mu_n(t) = 1-q_n, \quad s_n \xrightarrow{\tau} \delta_t, \quad t \xrightarrow{a} \delta_u$$

which is depicted in Fig. 4. The states s_n, $n \in \mathbb{N}$, and t are pairwise x-bisimilar for any $x \in \{b, \eta, w, d\}$, and none of them is R-divergent for the equivalence relation R with the two

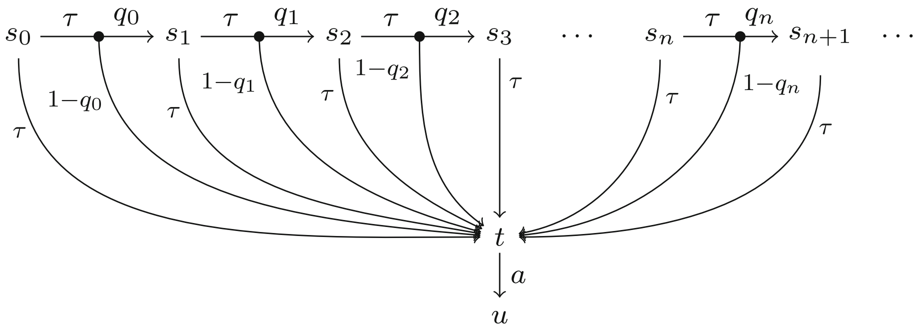

Fig. 4 Positive divergence in infinite MDP

equivalence classes $C_1 = \{t\} \cup \{s_n : n \in \mathbb{N}\}$ and $C_2 = \{u\}$. Note that under each scheduler the probability to reach t from s_0 is at least $1-q$. Thus, there is no scheduler that generates paths in $(C_1 \times \{\tau\})^\omega$ with probability 1. This explains why none of the states is R-divergent. In particular, R is ∇-respecting.

The R-divergence probability values are as follows: $\mathbb{DP}_\nabla(t, R) = \{0\}$ and $\mathbb{DP}_\nabla(s_n, R) = \{0, p_n\}$ where $p_n = q/q'_n$ with $q'_n = q_0 \cdot q_1 \cdot \ldots \cdot q_{n-1}$, i.e., $p_n = \prod_{j \geqslant n} q_n$. Then, the p_n's are in $]0, 1[$ and converge to 1.

Hence, R is a ∇-respecting bisimulation, while $\sup \mathbb{DP}_\nabla(s_0, R) = q > 0$. Thus, this example illustrates that statements (a) and (c) in Lemma 10 are no longer equivalent when dealing with countable state spaces. □

Lemma 11 (∇-respecting bisimulations and divergence probabilities) *Let R be a x-bisimulation where $x \in \{b, \eta, w, d\}$ in a finite MDP. Then, the following statements are equivalent:*

(a) *R is ∇-respecting.*
(b) *Whenever $(s, t) \in R$ then $\sup \mathbb{DP}_\nabla(s, R) = \sup \mathbb{DP}_\nabla(t, R)$.*
(c) *Whenever $(s, t) \in R$ then $\mathbb{DP}_\nabla(s, R) = \mathbb{DP}_\nabla(t, R)$.*

Proof The implication (c) $\Longrightarrow$ (b) is trivial.

To prove (b) $\Longrightarrow$ (a), suppose $\nabla^R(s)$ and $(s, t) \in R$. Then:

$$\sup \mathbb{DP}_\nabla(t, R) = \sup \mathbb{DP}_\nabla(s, R) = 1$$

As stated in Lemma 9, there is a scheduler achieving the maximal divergence probability. Thus, $\nabla^R(t)$.

The most interesting part is the proof of the implication (a) $\Longrightarrow$ (c). So, suppose R is ∇-respecting and $(s, t) \in R$. By symmetry it suffices to show that $\mathbb{DP}_\nabla(s, R) \subseteq \mathbb{DP}_\nabla(t, R)$. Pick an arbitrary scheduler σ. The task is to prove that there is a scheduler σ' such that

$$q_C \stackrel{\text{def}}{=} \Pr_s^\sigma\big((C \times \{\tau\})^\omega\big) = \Pr_t^{\sigma'}\big((C \times \{\tau\})^\omega\big)$$

where $C = [s]_R = [t]_R$.

For each R-equivalence class $D \in S/R$ with $D \neq C$, let

$$q_D = \Pr_s^\sigma\big(C_\tau \mathrm{U}\, D\big)$$

where $C_\tau \mathrm{U}\, D$ stands for the set of maximal paths that have a prefix of the form $s_0\, \tau\, s_1 \tau \ldots \tau s_n\, \tau\, u$ where $\{s_0, \ldots, s_n\} \subseteq C$ and $u \in D$. Moreover, we define

$$q_{vis} = \Pr_s^\sigma\big(C_\tau \mathrm{U}_{vis}\, S\big)$$

where $C_\tau \mathrm{U}_{vis}\, S$ denotes the set of maximal paths that have a prefix of the form $s_0\, \tau\, s_1 \tau \ldots \tau s_n\, a\, u$ where $\{s_0, \ldots, s_n\} \subseteq C$ and a is a visible action (i.e., $a \in \mathrm{Act}\backslash\{\tau\}$). Then,

$$q_{vis} + \sum_{D \in S/R} q_D = 1$$

and σ induces a weak transition $s \Rightarrow_c \mu$ where $\mu(D) = q_D$ for all $D \in S/R$, $D \neq C$ and $\mu(C) = q_C + q_{vis}$. By Lemma 2 there is a weak transition $t \Rightarrow_c \nu$ with $\mu \equiv_R \nu$. Let $\mathcal{T}$ be an associated compressed τ-tree for this weak transition from t.

If $q_C > 0$ then there is an end component $\mathcal{E}$ that consists of C-states and τ-transitions. But then the states in $\mathcal{E}$ are R-divergent. As R is ∇-respecting, this yields that all states in C are R-divergent. In particular, t is R-divergent. Let σ_C be a memoryless scheduler witnessing the R-divergence of all C-states.

Let us first consider the case $q_{vis} = 0$. We now consider any scheduler σ', which when started in t, realizes the weak transition $t \Rightarrow_c \nu$ by following the decisions in $\mathcal{T}$. As soon as a leaf v of $\mathcal{T}$ has been reached then σ' behaves as σ_C from then on. Note that if $state(v) \in C$ for some leaf v of $\mathcal{T}$ then $q_C > 0$, in which case all C-states are R-diverging (see above). Hence, we then have

$$\Pr_t^{\sigma'}\big((C \times \{\tau\})^\omega \big) = q_C,$$

which yields the claim.

The argument for $q_{vis} > 0$ is as follows. If q_{vis} is positive, then there is at least one state $s' \in C$ that has a transition $s' \xrightarrow{a} \theta$ for some visible action a. But then for all states $u \in C$ $u \Rightarrow_c \xrightarrow{a}_c \theta_u$ for some distribution θ_u. We now consider any (randomized) scheduler σ', which when started in t, realizes the weak transition $t \Rightarrow_c \nu$ by following the decisions in $\mathcal{T}$.

- As soon as a leaf v of $\mathcal{T}$ with $state(v) \notin C$ has been reached then σ' behaves as σ_C from then on.
- As soon as a leaf v of $\mathcal{T}$ with $state(v) \in C$ has been reached (this happens with probability $q_{vis}+q_C$), then σ' realizes $state(v) \Rightarrow_c \xrightarrow{a}_c \theta_{state(v)}$ with probability $q_{vis}/(q_{vis}+q_C)$ and behaves in an arbitary way after having performed the visible action a. With the remaining probability $q_C/(q_{vis}+q_C)$, σ' behaves as σ_C.

But then $\Pr_t^{\sigma'}((C \times \{\tau\})^\omega) = q_C$. □

Remark 2 (Positive explicit divergence probability) Given an equivalence R on the state space S, let us define the predicate $\nabla_{>0}^R$ by:

$$\nabla_{>0}^R(s) \quad \text{iff} \quad \exists \sigma \text{ s.t. } \Pr_s^\sigma\big(([s]_R \times \{\tau\})^\omega \big) > 0$$

which is obviously equivalent to the statement $\sup \mathbb{DP}_\nabla(s, R) > 0$. Given an x-bisimulation R, we say R is $\nabla_{>0}$-respecting if and only if the following condition (*) holds:

$$(s,t) \in R \text{ and } \nabla_{>0}^R(s) \text{ implies } \nabla_{>0}^R(t) \tag{*}$$

By Lemma 10, each ∇-respecting x-bisimulation is $\nabla_{>0}$-respecting. To see this, suppose R is an ∇-respecting x-bisimulation. Let $(s,t) \in R$ and $\nabla_{>0}^R(s)$. By the equivalence of statements (a) and (c) in Lemma 10 we obtain $\nabla^R(s)$. As R is ∇-respecting and s and t are R-equivalent we obtain $\nabla^R(t)$. But then $\nabla_{>0}^R(t)$. Thus, R satisfies condition (*).

The converse, however, does not hold. That is, there are $\nabla_{>0}$-respecting x-bisimulations that are not ∇-respecting.

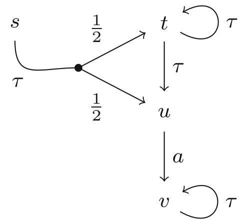

Fig. 5 $s \approx_x^{\nabla_{>0}} t$ but $s \not\approx x\nabla t$

Consider the MDP depicted in Fig. 5 consisting of states s, t, u, v with the transitions

$$\begin{array}{llll} s \xrightarrow{\tau} \frac{1}{2}\delta_u + \frac{1}{2}\delta_t & & & \\ t \xrightarrow{\tau} \delta_t & & t \xrightarrow{\tau} \delta_u & \\ u \xrightarrow{a} \delta_v & & v \xrightarrow{\tau} \delta_v & \end{array}$$

Then, s and t are x-bisimilar for any $x \in \{b, \eta, w, d\}$. State t is R-diverging for each x-bisimulation R, while state s is not. Thus, $s \not\approx x\nabla t$. On the other hand, if R is the equivalence that identifies s and t, but no other states, then R is a $\nabla_{>0}$-respecting x-bisimulation as both s and t have positive R-divergence probability under the memoryless scheduler σ that schedules the τ-labeled self-loop at state t. Note that the R-divergence probability for t under σ equals 1, while the R-divergence probability for s under σ is 1/2.

Thus, ∇-respecting x-bisimilarity (equivalence $\approx_x^{\nabla}$) is strictly finer than $\nabla_{>0}$-respecting x-bisimilarity (equivalence $\approx_x^{\nabla_{>0}}$) in finite MDPs.

In the case of (countably) infinite MDPs, both equivalences are incomparable. Indeed, take the example of Remark 1. As already pointed out, there $s_0 \approx_x^{\nabla} t$. However, $\nabla_{>0}(s_0)$ but $\neg\nabla_{>0}(t)$ and hence $s_0 \not\approx x\nabla_{>0}t$. □

5.2 Δ-divergence

Analogous results can be established for the divergence operator Δ. Given a state s and a scheduler σ. the value $\Pr_s^\sigma\big((S \times \{\tau\})^\omega\big)$ is called the Δ-divergence probability of state s under scheduler σ.

The following lemma holds in arbitrary (possibly countable) MDPs:

Lemma 12 (MD-schedulers witnessing Δ-divergence) *The set $S_\Delta = \{s \in S \mid \Delta(s)\}$ consisting of all Δ-divergent states is the largest set $W \subseteq S$ such that every state $s \in W$ has a transition $s \xrightarrow{\tau} \nu_s$ with $supp(\nu_s) \subseteq W$. In particular, there is a memoryless deterministic scheduler σ_Δ such that for each state $s \in S$:*

$$\Delta(s) \text{ iff } \Pr_s^{\sigma_\Delta}\big((S \times \{\tau\})^\omega\big) = 1$$

Proof Clearly, each state $s \in S_\Delta$ has a τ-transition $s \xrightarrow{\tau} \nu_s$ with $\text{supp}(\nu_s) \subseteq S_\Delta$. Vice verca, whenever $W \subseteq S$ such that every state $s \in W$ has a transition $s \xrightarrow{\tau} \nu_s$ with $\text{supp}(\nu) \subseteq W$ then for each memoryless deterministic scheduler σ that schedules $s \xrightarrow{\tau} \nu_s$ for each $s \in W$ and an arbitrary transition for all states in $S \backslash W$, we have $\Pr_s^\sigma\big((W \times \{\tau\})^\omega\big) = 1$ for all states $s \in W$. Hence, $W \subseteq S_\Delta$. □

We define $\mathbb{DP}_\Delta(s)$ as the set of all Δ-divergence probabilities of state s when ranging over all schedulers. That is:

$$\mathbb{DP}_\Delta(s) = \left\{ \Pr_s^\sigma\big((S \times \{\tau\})^\omega \big) : \sigma \text{ is a scheduler for } \mathcal{M} \right\}$$

Then, $\mathbb{DP}_\Delta(s)$ is a subset of $[0, 1]$. Let us first observe that extremal Δ-divergence probabilities are achieved by memoryless deterministic schedulers. To establish this result we rely on the default assumption that the given MDP $\mathcal{M}$ is finite.

Lemma 13 (MD-scheduler with extremal Δ-divergence probabilities) *In finite MDPs, there exist memoryless deterministic schedulers $\sigma_\Delta^{\max}$ and $\sigma_\Delta^{\min}$ such that*

$$\Pr_u^{\sigma_\Delta^{\max}}\big((S \times \{\tau\})^\omega \big) = \sup \mathbb{DP}_\Delta(u)$$
$$\Pr_u^{\sigma_\Delta^{\min}}\big((S \times \{\tau\})^\omega \big) = \inf \mathbb{DP}_\Delta(u)$$

for all states $u \in S$.

Proof Consider the sub-MDP $\mathcal{M}_\tau$ of $\mathcal{M}$ with the same state space S where $\rightarrow$ is restricted to $\mathcal{M}$'s τ-transitions. Let *Term* denote the set of terminal states in $\mathcal{M}_\tau$. Using standard results for finite-state MDPs, there is a memoryless deterministic scheduler σ^+ that minimizes the probability of reaching *Term* from every state. It is easy to see that any MD-extension $\sigma_\Delta^{\max}$ of σ^+ to a scheduler for $\mathcal{M}$ maximizes the probability for generating infinite τ-paths from every state.

The argument for $\sigma_\Delta^{\min}$ is analogous, but here we consider the set *Vis* of states $u \in S$ that have at least one visible transition $u \xrightarrow{a} \theta$. We then pick a memoryless deterministic scheduler σ^- for $\mathcal{M}_\tau$ that maximizes the probability to reach a state in $\mathit{Vis} \cup \mathit{Term}$ from every state and consider any memoryless deterministic scheduler $\sigma_\Delta^{\min}$ that behaves as σ^- for the states $u \in S \setminus \mathit{Vis}$ and that schedules some visible transition for the states in *Vis*. It is easy to see that this scheduler achieves the minimal Δ-divergence probabilities. $\square$

As a consequence of Lemma 13 we obtain:

$$\Delta(s) \quad iff \quad \max \mathbb{DP}_\Delta(s) = 1 \quad iff \quad \sup \mathbb{DP}_\Delta(s) = 1$$

which again is equivalent to the existence of a scheduler σ such that all maximal σ-paths from s belong to $(S \times \{\tau\})^\omega$. Thus, the analogous statements (a), (b), (d) in Lemma 10 rephrased for Δ (rather than ∇) are equivalent too. However, in contrast to Lemma 10, the statement $\sup \mathbb{DP}_\Delta(s) > 0$ does not imply $\Delta(s)$. A simple example for this phenomenon is the MDP of Remark 2 with $R = \text{id}$ (the identity relation). Obviously, R is a Δ-respecting x-bisimulation for any $x \in \{b, \eta, w, d\}$ and $\sup \mathbb{DP}_\Delta(s) = 1/2$. Thus, $\sup \mathbb{DP}_\Delta(s) > 0$, while $\neg\Delta(s)$.

Nevertheless, the equivalence of the three statements in Lemma 11 holds in an analogous way. For establishing this result (see Lemma 14 below), we shall rely on Lemma 13 which requires the finiteness of the given MDP.

Lemma 14 (Δ-respecting bisimulations and Δ-divergence probabilities) *Let R be a x-bisimulation where $x \in \{b, \eta, w, d\}$ in a finite MDP. Then, the following statements are equivalent:*

(a) *R is Δ-respecting.*
(b) *Whenever $(s, t) \in R$ then* $\sup \mathbb{DP}_\Delta(s) = \sup \mathbb{DP}_\Delta(t)$.
(c) *Whenever $(s, t) \in R$ then* $\mathbb{DP}_\Delta(s) = \mathbb{DP}_\Delta(t)$.

Proof The implication (c) $\Longrightarrow$ (b) is trivial.

To prove (b) $\Longrightarrow$ (a), suppose $\Delta(s)$ and $(s,t) \in R$. Then:

$$\sup \mathbb{DP}_\Delta(t) = \sup \mathbb{DP}_\Delta(s) = 1$$

where the first equality holds by assumption (b) and the second equality as we assume $\Delta(s)$. By Lemma 13, we have $\sup \mathbb{DP}_\Delta(u) = \max \mathbb{DP}_\Delta(u)$ for all states u. In particular, $\max \mathbb{DP}_\Delta(t) = 1$, which yields $\Delta(t)$.

To prove (a) $\Longrightarrow$ (c), we suppose that R is a Δ-respecting x-bisimulation and pick a pair $(s,t) \in R$. By symmetry it suffices to show that $\mathbb{DP}_\Delta(s) \subseteq \mathbb{DP}_\Delta(t)$. Pick an arbitrary scheduler σ. The task is to prove that there is a scheduler σ' such that:

$$\Pr_s^\sigma\big((S \times \{\tau\})^\omega\big) = \Pr_t^{\sigma'}\big((S \times \{\tau\})^\omega\big)$$

We will first treat the case of schedulers that are memoryless and deterministic (see Claim 1 below), and then show that the Δ-divergence probability for state s under the given (possibly randomized and not memoryless) scheduler σ is achieved by some scheduler which can be seen as probabilistic choice of two memoryless deterministic schedulers (Claim 2). Finally, we combine the results of Claim 1 and Claim 2 to derive the existence of a scheduler σ' with the desired property (Claim 3).

Claim 1. If σ_0 is memoryless deterministic then there is a scheduler σ_0' with

$$\Pr_s^{\sigma_0}\big((S \times \{\tau\})^\omega\big) = Pr_t^{\sigma_0'}\big((S \times \{\tau\})^\omega\big)$$

Proof of Claim 1. As σ_0 is supposed to be memoryless and deterministic, the behavior of $\mathcal{M}$ under σ_0 can be represented by a finite Markov chain $\mathcal{C}$ with state space S and where the transition probabilities of each state u in $\mathcal{C}$ are given by the transition scheduled by σ_0 for state u. That is, if σ_0 schedules the transition $u \xrightarrow{\alpha_u} \theta_u$ for state u then $P_\mathcal{C}(u,u') = \theta_u(u')$.

Let T denote the set of states $u \in S$ where σ_0 schedules a τ-transition, i.e., where $\alpha_u = \tau$. Thus, σ_0 schedules transitions with a visible action for the states in $Vis \stackrel{\text{def}}{=} S \backslash T$. Let E denote the set of states that belong to a bottom strongly connected component (BSCC) of $\mathcal{C}$ consisting of states in T.[3] Note that these BSCCs of $\mathcal{C}$ correspond to end components of $\mathcal{M}$ that are built by τ-actions, briefly called τ-ECs. Then,

$$q_\Delta \stackrel{\text{def}}{=} \Pr_s^{\sigma_0}\big((S \times \{\tau\})^\omega\big) = Pr_s^{\sigma_0}(T \,\mathrm{U}\, E)$$

where U denotes the LTL until operator. That is, if $A, B \subseteq S$ then $A \,\mathrm{U}\, B$ denotes the set of paths that start with a (possibly empty) prefix of A-states followed by a B-state. With $q_{vis} = \Pr_s^{\sigma_0}(T \,\mathrm{U}\, Vis)$ we have $q_\Delta + q_{vis} = 1$.

Let us write $[Vis]_R$ and $[E]_R$ for the R-closure of Vis and E, respectively. That is, $u \in [E]_R$ iff there is some $u' \in E$ such that u and u' are R-equivalent, and the analogous statement for $[Vis]_R$.

We shall use the following ingredients to define scheduler σ_0':

– The given scheduler σ_0 induces a weak transition $s \Rightarrow_c \mu$ such that $\mu(u) = \Pr_s^{\sigma_0}(T \,\mathrm{U}\, u)$ for all states $u \in Vis \cup E$ and $\mu(u) = 0$ for all other states. As s and t are R-equivalent and R is an x-bisimulation there is a weak transition $t \Rightarrow_c \nu$ with $\mu \equiv_R \nu$ (Lemma 2). Let $\mathcal{T}$ be a corresponding compressed τ-tree for $t \Rightarrow_c \nu$.

[3] If there is no BSCC in $\mathcal{C}$ that consists of T-states then $E = \varnothing$, in which case $\Pr_u^{\sigma_0}\big((S \times \{\tau\})^\omega\big) = 0$ for all states $u \in S$.

- For each R-equivalence class $D \in S/R$ where $D \cap \mathit{Vis}$ is nonempty and each state $u \in D$ we pick a visible action a_u such that $u \Rightarrow_c \xrightarrow{a_u}$. (If $u \in D \cap \mathit{Vis}$ then we can deal with $a_u = \alpha_u$. For states $u \in D \backslash \mathit{Vis}$ we can pick an arbitrary state $u_0 \in D \cap \mathit{Vis}$ and deal with $a_u = \alpha_{u_0}$.)
- As $\Delta(u)$ for all states $u \in E$ and R is Δ-respecting, we have $\Delta(u)$ for all states $u \in [E]_R$. Let σ_Δ be a memoryless deterministic scheduler witnessing $\Delta(u)$ for all states $u \in [E]_R$. For example, we may deal with $\sigma_\Delta = \sigma_\Delta^{\max}$ as in Lemma 13.

Let now σ_0' be any scheduler that, when started in state t, first follows the decisions in the τ-tree $\mathcal{T}$ for the $t \Rightarrow_c \nu$. Having reached a leaf v of $\mathcal{T}$, the behavior of σ_0' depends on the state $u \stackrel{\text{def}}{=} \mathit{state}(v)$. Note that $u \in \mathrm{supp}(\nu)$, and hence, u is R-equivalent to some state in $\mathrm{supp}(\mu)$.

Let us first consider the case $\mu([\mathit{Vis}]_R \cap [E]_R) = 0$, in which case $u \in S \backslash ([\mathit{Vis}]_R \cap [E]_R)$. In this case, we define the behavior of σ_0' after having reached u as follows:

- If $u \in [\mathit{Vis}]_R$ then scheduler σ_0' mimicks $u \Rightarrow_c \xrightarrow{a_u}$ and behaves in an arbitrary way after having performed the visible action a_u.
- If $u \in [E]_R$ then scheduler σ_0' behaves as σ_Δ from then on.

As $\mu \equiv_R \nu$ and $\mathrm{supp}(\mu) \subseteq \mathit{Vis} \cup E$ we have $\mathrm{supp}(\nu) \subseteq [\mathit{Vis}]_R \cup [E]_R$ and

$$\Pr_t^{\sigma_0'}\big((S \times \{\tau\})^\omega\big) \;=\; \nu([E]_R) = mu(E) = \; q_\Delta$$

Suppose now that $\mu([\mathit{Vis}]_R \cap [E]_R)$ is positive. The states $u \in [\mathit{Vis}]_R \cap [E]_R$ satisfy $\Delta(u)$, but they can also perform visible actions after some τ's. Let p_{vis} denote the probability under σ_0 started in state s to perform a visible action in some state $u \in [E]_R$ after a sequence of τ's taken in T-states, and let p_{div} the probability to enter a state $u \in [\mathit{Vis}]_R \cap E$ from some state in $T \backslash E$. Formally:

$$\begin{aligned} p_{vis} &= \Pr_s^{\sigma_0}\big((T \backslash E)\,\mathrm{U}\,(\mathit{Vis} \cap [E]_R)\big) = \Pr_s^{\sigma_0}\big(T\,\mathrm{U}\,(\mathit{Vis} \cap [E]_R)\big) \\ p_{div} &= \Pr_s^{\sigma_0}\big((T \backslash E)\,\mathrm{U}\,([\mathit{Vis}]_R \cap E)\big) \end{aligned}$$

Note that $\mathit{Vis} \cap E = \varnothing$ and therefore the sets $\mathit{Vis} \cap [E]_R$ and $[\mathit{Vis}]_R \cap E$ are disjoint. Moreover, let

$$p = \frac{p_{div}}{p_{vis} + p_{div}}$$

denote the conditional Δ-divergence probability for s under σ_0 under the condition that σ_0 enters a state in $[\mathit{Vis}]_R \cap [E]_R$ from some state in $T \backslash E$. Note that:

$$p_{vis} + p_{div} = \mu([\mathit{Vis}]_R \cap [E]_R) \geqslant \mu(\mathit{Vis} \cap E) > 0$$

Then, we define the behavior of σ_0' after having reached a leaf v in $\mathcal{T}$ with $\mathit{state}(v) = u$ as follows:

- If $u \in [\mathit{Vis}]_R \backslash [E]_R$ then σ_0' mimicks $u \Rightarrow_c \xrightarrow{a_u}$ and behaves in an arbitrary way after having performed the visible action a_u.
- If $u \in [E]_R \backslash [\mathit{Vis}]_R$ then σ_0' behaves as σ_Δ from then on.
- If $u \in [E]_R \cap [\mathit{Vis}]_R$ then σ_0' behaves randomized:
  - with probability p, scheduler σ_0' behaves as σ_Δ from then on,
  - with probability $1-p$, scheduler σ_0' mimicks $u \Rightarrow_c \xrightarrow{a_u}$ with arbitrary behavior afterwards.

Thus, the Δ-divergence probability of t under σ_0' is:

$$\Pr_t^{\sigma_0'}\big((S \times \{\tau\})^\omega\big) = \nu([E]_R \backslash [Vis]_R) \ + \ p \cdot \nu([Vis]_R \cap [E]_R)$$

As $[Vis]_R$ and $[E]_R$ are R-closed, so are the sets $[Vis]_R \backslash [E]_R$, $[E]_R \backslash [Vis]_R$ and $[Vis]_R \cap [E]_R$. Hence, $\mu \equiv_R \nu$ yields:

$$\begin{aligned} \mu([Vis]_R \backslash [E]_R) &= \nu([Vis]_R \backslash [E]_R) \\ \mu([E]_R \backslash [Vis]_R) &= \nu([E]_R \backslash [Vis]_R) \\ \mu([Vis]_R \cap [E]_R) &= \nu([Vis]_R \cap [E]_R) \end{aligned}$$

As $\mu(Vis \cup E) = 1$ and $Vis \cap E = \varnothing$ we have:

$$\mu(E \backslash [Vis]_R) = \mu([E]_R \backslash [Vis]_R) \text{ and } \mu(Vis \backslash [E]_R) = \mu([Vis]_R \backslash [E]_R)$$

Putting things together, we obtain:

$$\begin{aligned} q_\Delta &= \mu(E \backslash [Vis]_R) + p_{div} \\ &= \mu([E]_R \backslash [Vis]_R) + p \cdot \mu([Vis]_R \cap [E]_R) \\ &= \nu([E]_R \backslash [Vis]_R) + p \cdot \nu([Vis]_R \cap [E]_R) \\ &= \Pr_t^{\sigma_0'}\big((S \times \{\tau\})^\omega\big) \end{aligned}$$

This completes the proof of Claim 1.

Notation: probabilistic choice of schedulers. Given two schedulers σ_1, σ_2 and $p \in [0, 1]$, let $\sigma_1 \oplus_p \sigma_2$ denotes the scheduler which initially tosses a biased coin that yields head with probability p and tail with probability $1-p$. If the outcome of the coin tossing is head then $\sigma_1 \oplus_p \sigma_2$ behaves as σ_1 from then on. If the outcome of the coin tossing is tail then $\sigma_1 \oplus_p \sigma_2$ behaves as σ_2 from then on.

Claim 2. There exist memoryless deterministic schedulers σ_1 and σ_2 and $p \in [0, 1]$ such that

$$\Pr_s^{\sigma''}\big((S \times \{\tau\})^\omega\big) = \Pr_s^{\sigma}\big((S \times \{\tau\})^\omega\big)$$

where $\sigma'' = \sigma_1 \oplus_p \sigma_2$.

Proof of Claim 2. Obviously, $\min \mathbb{DP}_\Delta(s) \leqslant q_\Delta \leqslant \max \mathbb{DP}_\Delta(s)$. Hence, there is a real number $p \in [0, 1]$ such that:

$$\Pr_s^{\sigma}\big((S \times \{\tau\})^\omega\big) = p \cdot \min \mathbb{DP}_\Delta(s) + (1-p) \cdot \max \mathbb{DP}_\Delta(s)$$

The statement of Claim 2 follows by considering $\sigma_1 = \sigma_\Delta^{\min}$ and $\sigma_2 = \sigma_\Delta^{\max}$ where $\sigma_\Delta^{\min}$ and $\sigma_\Delta^{\max}$ are as in Lemma 13.

Claim 3. There exists a scheduler σ' with $\Pr_s^{\sigma}\big((S \times \{\tau\})^\omega\big) = \Pr_t^{\sigma'}\big((S \times \{\tau\})^\omega\big)$.

Proof of Claim 3. Let σ_1, σ_2 and p be as in Claim 2. We can now rely on Claim 1 to obtain schedulers σ_1' and σ_2' such

$$\Pr_s^{\sigma_i}\big((S \times \{\tau\})^\omega\big) = Pr_t^{\sigma_i'}\big((S \times \{\tau\})^\omega\big)$$

for $i = 1, 2$. The scheduler $\sigma' = \sigma_1' \oplus_p \sigma_2'$ then fulfills the desired property. □

Remark 3 (Positive Δ-divergence probability) Analogous to the considerations in Remark 2 we introduce the predicate $\Delta_{>0}$ by:

$$\Delta_{>0}(s) \quad \text{iff} \quad \exists \sigma \text{ s.t. } \Pr_s^{\sigma}\big((S \times \{\tau\})^\omega\big) > 0$$

Clearly, we then have $\Delta_{>0}(s)$ if and only if $\sup \mathbb{DP}_{\Delta}(s) > 0$. Given an x-bisimulation R, we say R is $\Delta_{>0}$-respecting if and only if the following condition (*) holds:

$$(s,t) \in R \text{ and } \Delta_{>0}(s) \text{ implies } \Delta_{>0}(t)$$

In Remark 2 we saw that ∇-respecting x-bisimilarity is strictly finer than $\nabla_{>0}$-respecting x-bisimilarity. To prove the "is-finer-than" relationship we used the equivalence of statements (a) and (c) in Lemma 10. Although the analogous statements are not equivalent for the Δ-predicate and Δ-divergence probabilities instead of the ∇-predicate and explicit divergence probabilities (see above), Δ-respecting x-bisimilarity $\approx_x^{\Delta}$ is strictly finer than $\Delta_{>0}$-respecting x-bisimilarity $\approx_x^{\Delta_{>0}}$ in finite-state MDPs. Formally, for each $x \in \{b, \eta, w, d\}$, the following statement holds in finite MDPs:

$$\text{Each } \Delta\text{-respecting } x\text{-bisimulation is } \Delta_{>0}\text{-respecting.} \tag{†}$$

To prove statement (†), consider a Δ-respecting x-bisimulation R and a state-pair $(s,t) \in R$ with $\Delta_{>0}(s)$. The goal is to show $\Delta_{>0}(t)$. For this, we consider the memoryless deterministic scheduler $\sigma = \sigma_{\Delta}^{\max}$ of Lemma 13. Then, $\Pr_s^{\sigma}\big((S \times \{\tau\})^{\omega}\big) > 0$. Let $\mathcal{C}$ be the finite-state Markov chain with state space S associated by σ. Then, there is BSCC $\mathcal{B}$ of $\mathcal{C}$ that is a reachable from s via a τ-path and where E is built by τ-transitions of $\mathcal{M}$. (I.e., σ schedules a τ-transition for all states in that BSCC $\mathcal{B}$ of $\mathcal{C}$.) Clearly, for the states $u \in E$ we have $\Pr_u^{\sigma}\big((S \times \{\tau\})^{\omega}\big) = 1$. Thus, $\Delta(u)$ for all states $u \in \mathcal{B}$.

The simple case is where $\mathcal{B} \cap [s]_R$ is nonempty as then we can pick a state $u \in \mathcal{B}$ that is R-equivalent to s and t. As R is Δ-respecting and $\Delta(u)$ we get $\Delta(s)$ and $\Delta(t)$. The latter implies $\Delta_{>0}(t)$. The argument for the general case (no matter whether $\mathcal{B} \cap [s]_R$ is empty or not) is as follows. Scheduler σ induces a weak transition $s \Rightarrow_c \mu$ where $\mathrm{supp}(\mu) \cap \mathcal{B} \neq \varnothing$. By Lemma 2, there is a weak transition $t \Rightarrow_c \nu$ with $\nu \equiv_R \mu$. In particular, $\mathrm{supp}(\nu)$ contains a state u that is R-equivalent to some $\mathcal{B}$-state u'. As $\Delta(u')$ and R is Δ-respecting we obtain $\Delta(u)$, and therefore $\Pr_u^{\sigma}\big((S \times \{\tau\})^{\omega}\big) = 1$. (Recall that $\sigma = \sigma_{\Delta}^{\max}$.) But then we can consider any scheduler σ' that first realizes the weak transition $t \Rightarrow_c \nu$ by following the decisions of a corresponding compressed τ-tree $\mathcal{T}$ and that behaves as σ as soon as a leaf of $\mathcal{T}$ has been reached. As $\mathcal{T}$ has a leaf v with $state(v) = u$, this yields the existence of a σ'-path from t to u built by τ-transitions. Therefore, $\Pr_t^{\sigma'}\big((S \times \{\tau\})^{\omega}\big) > 0$. Hence, $\Delta_{>0}(t)$.

This completes the proof of statement (†), which yields that $\approx_x^{\Delta}$ is finer than $\approx_x^{\Delta_{>0}}$. The example provided in Remark 2 also illustrates that $\approx_x^{\Delta}$ is strictly finer than $\approx_x^{\Delta_{>0}}$, as we have $s \approx_x^{\Delta_{>0}} t$, but $s \not\approx x\Delta t$.

The above proof for statement (†) uses the finiteness of the given MDP. Indeed, statement (†) is wrong for countable (infinite) MDPs. To see this consider the MDP of Remark 1. In this MDP, the predicate Δ is empty, while $\Delta_{>0} = \{s_n : n \in \mathbb{N}\}$. Let R be the equivalence with the two equivalence classes $C_1 = \{t\} \cup \{s_n : n \in \mathbb{N}\}$ and $C_2 = \{u\}$. Then, R is a Δ-respecting x-bisimulation, but not $\Delta_{>0}$-respecting as $\Delta_{>0}(s_n)$ and $\Delta_{>0}(t)$. □

Remark 4 (Schedulers witnessing almost-sure resp. positive divergence) For the almost-sure divergence predicates Δ we established the existence of memoryless deterministic schedulers witnessing the Δ-divergence for all Δ-divergent states in arbitrary (possibly infinite) MDPs (see Lemma 12). The analogous result holds for $\Delta_{>0}$-divergence in finite MDPs, but not in infinite MDPs. An example for the latter is obtained by considering the countable MDP $\mathcal{M}$ with state space $S = \{t, u\} \cup \{s_n : n \in \mathbb{N}\}$, action set $\mathrm{Act} = \{a, \tau\}$ and the following transition relation. We pick a strictly increasing sequence $(q_n)_{n \in \mathbb{N}}$ of rational numbers q_n in the open interval $]0, 1[$ such that their product converges to a positive value $p \in]0, 1[$. For

$n \in \mathbb{N}$, let $\mu_n \in \mathrm{Dist}(S)$ be given by

$$\mu_n(s_0) = \mu_n(s_{n+1}) = \frac{1}{2} \cdot q_n \text{ and } \mu_n(t) = 1-q_n.$$

Then, the transition relation in $\mathcal{M}$ is as follows:

- State s_n for $n \geqslant 1$ has a single outgoing transition, namely $s_n \xrightarrow{\tau} \mu_n$.
- The outgoing transitions of state s_0 are $s_0 \xrightarrow{\tau} \mu_n$ for each $n \in \mathbb{N}$.
- State u is terminal and state t has a single transition, namely $t \xrightarrow{a} \delta_u$.

We now show that $\sup \mathbb{DP}_\Delta(s_0) = 1$, but there is no memoryless scheduler witnessing $\Delta_{>0}(s_0)$ and there is no scheduler achieving divergence probability 1 for state s_0. This is due to the following observations:

- For $k \in \mathbb{N}$, consider the deterministic scheduler σ_k that uses a step counter as memory and schedules $s_0 \xrightarrow{\tau} \mu_{k+\ell}$ if σ is in state s_0 after exactly ℓ steps. Then, all σ-paths from s_0 of length $\ell+1$ end in state u, t, s_0 or $s_{k+\ell+1}$. Thus, the probability for σ_k to be in one of the s-states after $\ell+1$ steps is $q_k \cdot q_{k+1} \ldots q_{k+\ell}$. But then, the Δ-divergence probability for s_0 under σ_k is $p_k \stackrel{\text{def}}{=} \prod_{n \geqslant k} q_k$. As the values p_k are strictly increasing and converge to 1, we obtain $\sup \mathbb{DP}_\Delta(s_0) = 1$.
- Under every scheduler, almost all Δ-divergent paths from s_0 visit s_0 infinitely often. If σ is memoryless, then σ moves from s_0 to t within one step with some positive probability ε. Thus, σ eventually moves to t with probability $\sum_{i=0}^{\infty} \varepsilon \cdot (1-\varepsilon)^i = 1$. This yields $\Pr_s^\sigma((S \times \{\tau\})^\omega) = 0$ for each memoryless scheduler σ.
- For each scheduler σ there is some positive probability to move from s_0 to t in the first step. This yields $\Pr_s^\sigma((S \times \{\tau\})^\omega) < 1$ for each scheduler σ.

To provide an example illustrating the analogous phenomenon for explicit divergence, we consider a variant $\mathcal{M}'$ of the above MDP $\mathcal{M}$. The state space and action set is the same as in $\mathcal{M}$. As in $\mathcal{M}$, state t has a single transition $t \xrightarrow{a} \delta_u$ and u is terminal. Let now $\nu_n \in \mathrm{Dist}(S)$ defined by:

$$\nu_n(s_0) = q_n^2,\ \nu_n(s_{n+1}) = (1-q_n) \cdot q_n \text{ and } \nu_n(t) = 1-q_n.$$

Then, for each $n \in \mathbb{N}$, the outgoing transitions of the state s_n are $s_n \xrightarrow{\tau} \nu_m$ for all $m \in \mathbb{N}$. Let now R be the equivalence relation on S with the three equivalence classes $\{s_n : n \in \mathbb{N}\}$, $\{t\}$ and $\{u\}$. As all s-states have the same transitions, R is an x-bisimulation for each $x \in \{b, \eta, w, d\}$. We then have:

- There is no scheduler σ where the R-divergence probability of state s_0 is 1. (Already after the first transition, each scheduler reaches t from s_0 with some positive probability.)
- There is a sequence $(\sigma_k)_{k \in \mathbb{N}}$ of schedulers such that $\sup \Pr_{s_0}^{\sigma_k}\big(([s_0]_R \times \{\tau\})^\omega\big) = 1$. To see this, we regard the deterministic schedulers σ_k that use a step counter as memory. Whenever σ_k is in one of the s-states, say s_n, after ℓ steps then σ_k schedules the transition $s_n \xrightarrow{\tau} \nu_{k+\ell}$. Thus, the probability for σ_k to stay in the equivalence class of the s-states is p_k (defined as above as the infinite product of the values q_n for $n \geqslant k$). The claim follows as the sequence $(p_k)_{k \in \mathbb{N}}$ is strictly increasing and converges to 1.
- $\Pr_{s_0}^\sigma\big(([s_0]_R \times \{\tau\})^\omega\big) = 0$ for each memoryless scheduler σ. To see this, we observe that for each scheduler the probability for generating an R-divergent path from s_0 that visits s_0 only finitely often is 0. This is because the values $(1-q_n)q_n$ are bounded by 1/2. Now, if σ is memoryless then there is some $\varepsilon > 0$, such that whenever σ visits s_0 then the probability to move from s_0 to t is ε. But since almost all R-divergent σ-paths visit s_0 infinitely often, the set of the R-divergent σ-paths must be a null set.

So, this example shows that in countable MDPs, $\sup \mathbb{DP}_{\nabla}(s, R) = 1$ is possible, although there is no scheduler that achieves probability 1 for the R-divergent paths from s and there is no memoryless scheduler witnessing $\nabla^R_{>0}(s)$.

Let us suppose now that $\mathcal{M}$ is a finite MDP. Then, there is a memoryless deterministic scheduler σ such that for all states $s \in S$:

$$\Delta_{>0}(s) \text{ iff } \mathrm{Pr}^{\sigma}_{s}\big((S \times \{\tau\})\big) > 0.$$

This result is a direct consequence of Lemma 13 as we may deal with $\sigma = \sigma^{\max}_{\Delta}$.

Likewise, if R is $\nabla_{>0}$-respecting x-bisimulation in a finite MDP, then there exist a memoryless deterministic scheduler σ such that

$$\nabla^R_{>0}(s) \text{ iff } \mathrm{Pr}^{\sigma}_{s}\big(([s]_R \times \{\tau\})\big) > 0.$$

Such a scheduler is obtained by considering a memoryless deterministic scheduler maximizing the R-divergence probability from every state (see Lemma 9). □

Remark 5 ($\Delta_{>0}$-respecting and $\nabla_{>0}$-respecting) After introducing the predicates $\nabla_{>0}$ and $\Delta_{>0}$, it also arises the question whether $\approx_x^{\nabla_{>0}}$ is finer than $\approx_x^{\Delta_{>0}}$. This actually holds, and we formally state it as follows. For each $x \in \{b, \eta, w, d\}$, the following statement (‡) holds in finite MDPs:

$$\text{Each stutter-closed } \nabla_{>0}\text{-respecting } x\text{-bisimulation is } \Delta_{>0}\text{-respecting.} \tag{‡}$$

Its proof follows the same logic as the proof of (†) above. Let R be a stutter-closed $\nabla_{>0}$-respecting x-bisimulation and let $(s, t) \in R$ with $\Delta_{>0}(s)$. The goal is to show $\Delta_{>0}(t)$. Consider the memoryless deterministic scheduler $\sigma = \sigma^{\max}_{\Delta}$ of Lemma 13. Then, $\mathrm{Pr}^{\sigma}_{s}\big((S \times \{\tau\})^{\omega}\big) > 0$. Let $\mathcal{C}$ be the finite-state Markov chain with state space S associated by σ. Then, there is BSCC $\mathcal{B}$ of $\mathcal{C}$ that is a reachable from s via a τ-path and where $\mathcal{B}$ is built by τ-transitions of $\mathcal{M}$. (I.e., σ schedules a τ-transition for all states in the BSCC $\mathcal{B}$ of $\mathcal{C}$.) Clearly, for every state $u \in \mathcal{B}$, $\mathrm{Pr}^{\sigma}_{u}\big((\mathcal{B} \times \{\tau\})^{\omega}\big) = 1$. Also, each state in $\mathcal{B}$ reaches any other state in $\mathcal{B}$ with probability 1. ($\mathcal{B}$ is a BSCC in the Markov chain $\mathcal{C}$.) Thererore, for every $u, u' \in \mathcal{B}$, we have $u \Rightarrow_c \delta_{u'}$ and $u' \Rightarrow_c \delta_u$. Since R is stutter-closed, by Corollary 3, $(u, u') \in R$, for all $u, u' \in \mathcal{B}$. Hence, for all $u \in \mathcal{B}$, we have $\mathrm{Pr}^{\sigma}_{u}\big(([u]_R \times \{\tau\})^{\omega}\big) = 1$ since $\mathcal{B} \subseteq [u]_R$. Therefore, $\nabla^R_{>0}(u)$.

Notice that σ induces a weak transition $s \Rightarrow_c \mu$ where $\mathrm{supp}(\mu) \cap \mathcal{B} \neq \varnothing$. By Lemma 2, there exists ν s.t. $t \Rightarrow_c \nu$ and $\nu \equiv_R \mu$. In particular, there must exists a state $u \in \mathrm{supp}(\nu)$ such that $(u', u) \in R$ for some $u' \in \mathrm{supp}(\mu) \cap \mathcal{B}$. As $\nabla_{>0}(u')$ and R is $\nabla_{>0}$-respecting, we get $\nabla_{>0}(u)$, and therefore $\mathrm{Pr}^{\sigma'}_{u}\big((S \times \{\tau\})^{\omega}\big) \geq \mathrm{Pr}^{\sigma'}_{u}\big(([u]_R \times \{\tau\})^{\omega}\big) > 0$. Construct the scheduler σ'' that first realizes the weak transition $t \Rightarrow_c \nu$ by following the decisions of a corresponding compressed τ-tree $\mathcal{T}$ and that behaves as σ' as soon as a leaf of $\mathcal{T}$ has been reached. As $\mathcal{T}$ has a leaf v with $\mathit{state}(v) = u$, this yields the existence of a σ''-path from t to u built by τ-transitions. Therefore, $\mathrm{Pr}^{\sigma'}_{t}\big((S \times \{\tau\})^{\omega}\big) > 0$. Hence, $\Delta_{>0}(t)$, which completes the proof of (‡). □

6 Modal logics for probabilities and silent moves

We are now turning to the question in how far the various relations we discussed thus far can be captured by modal logics. We indeed present logics with modalities along the two dimensions that altogether can be shown to characterise the probabilistic bisimulation spectrum with silent moves. Albeit written in a modal logic style, the concepts are similar in spirit to the canonical testers developed by van Glabbeek for the nonprobabilistic spectrum [23].

Syntactic ingredients For $x \in \{b, \eta, d, w\}$ and $\xi \in \{\nabla, \Delta, \mathsf{s}, 0, \lambda, \varnothing\}$, we define the logics $\mathcal{L}^{state}_{x,\xi}$ and $\mathcal{L}^{dist}_{x,\xi}$ containing all formulas generated by the following respective grammars:

$$\begin{aligned}
\mathcal{L}^{state}_{x,\xi}: \quad & \phi ::= \top \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid modact_x \mid moddiv^{\xi}_x \\
\mathcal{L}^{dist}_{x,\xi}: \quad & \psi ::= \neg\psi \mid \psi_1 \wedge \psi_2 \mid [\phi]_{\bowtie q}
\end{aligned}$$

where $q \in [0, 1]$, $\bowtie$ is a comparison operator in $\{<, \leq\}$ and $modact_x$ and $moddiv^{\xi}_x$ are as follows, with $\alpha \in \mathrm{Act}$:

x	b	η	d	w
$modact_x$	$\langle\psi_1\langle\alpha\rangle\psi_2\rangle$	$\langle\psi_1\langle\alpha\rangle\!\rangle\psi_2$	$\langle\!\langle\alpha\rangle\psi\rangle$	$\langle\!\langle\alpha\rangle\!\rangle\psi$

ξ	$\varnothing$	0	s	λ	Δ	∇	
x	any	any	any	any	any	b or η	d or w
$moddiv^{\xi}_x$	no operator defined	0	s	λ	Δ	$\nabla\phi$	$\nabla_{\epsilon}\phi$

Notice that in particular α may be τ in $modact_x$.

Disjunction and other boolean operators can be derived as usual. Lower bounds and equality in the probability operator can be defined as follows:

$$[\phi]_{>q} \overset{\text{def}}{=} \neg[\phi]_{\leq q} \qquad [\phi]_{\geq q} \overset{\text{def}}{=} \neg[\phi]_{<q} \qquad [\phi]_{=q} \overset{\text{def}}{=} [\phi]_{\leq q} \wedge [\phi]_{\geq q}$$

Semantics of the logics The semantics of the logical operators are defined as follows using satisfaction relations $\models_{\mathcal{M}}$ for the state and distribution formulas over the states, respectively over distributions on states in a given MDP $\mathcal{M}$. Furthermore, we define the satisfaction relations for state and distribution formulas by:

$$\begin{aligned}
Sat_{\mathcal{M}}(\phi) &= \{s \in S \mid s \models_{\mathcal{M}} \phi\} \\
Sat_{\mathcal{M}}(\psi) &= \{\mu \in \mathrm{Dist}(S) \mid \mu \models_{\mathcal{M}} \psi\}
\end{aligned}$$

If $\mathcal{M}$ is clear from the context, we omit the subscript $\mathcal{M}$ and simply write $\models$ and $Sat(\cdot)$ rather than $\models_{\mathcal{M}}$ and $Sat_{\mathcal{M}}$. We define

$$\begin{aligned}
s \models \langle\psi_1\langle\alpha\rangle\psi_2\rangle \quad &\text{iff} \quad \exists \mu_1, \mu_2 : s \Rightarrow_c \mu_1 \xrightarrow{\alpha}_c \mu_2, \mu_1 \models \psi_1, \text{ and } \mu_2 \models \psi_2 \\
s \models \langle\psi_1\langle\alpha\rangle\!\rangle\psi_2 \quad &\text{iff} \quad \exists \mu_1, \mu_2 : s \Rightarrow_c \mu_1 \xrightarrow{\alpha}_c \Rightarrow_c \mu_2, \mu_1 \models \psi_1, \text{ and } \mu_2 \models \psi_2 \\
s \models \langle\!\langle\alpha\rangle\psi\rangle \quad &\text{iff} \quad \exists \mu : s \Rightarrow_c \xrightarrow{\alpha}_c \mu \text{ and } \mu \models \psi \\
s \models \langle\!\langle\alpha\rangle\!\rangle\psi \quad &\text{iff} \quad \exists \mu : s \Rightarrow_c \xrightarrow{\alpha}_c \Rightarrow_c \mu \text{ and } \mu \models \psi
\end{aligned}$$

The semantics of negation and conjunction is as expected. The meanings of the divergence modalities are defined as follows.

If $\xi \in \{\Delta, \mathsf{s}, 0, \lambda\}$ then:

$$s \models \xi \quad \text{iff} \quad \exists \mu \in \mathrm{Dist}(S) \text{ s.t. } s \Rightarrow_c \mu \text{ and } \xi(\mu)$$

where $\xi(\mu)$ stands short for "$\xi(s)$ for all $s \in \mathrm{supp}(\mu)$". Note that the weak transition in the semantics of formula Δ is irrelevant as we have $s \models \Delta$ iff $\xi(s)$.

The semantics of the two logical operators for $\xi = \nabla$ is given by:

$$s \models \nabla\phi \quad \text{iff} \quad \exists\sigma \text{ s.t. } \Pr_s^\sigma\left((Sat(\phi)\times\{\tau\})^\omega \right) = 1$$
$$s \models \nabla_\epsilon\phi \quad \text{iff} \quad \exists\sigma \text{ s.t. } \Pr_s^\sigma\left((U_\phi \times \{\tau\})^\omega\right) = 1$$

where U_ϕ denotes the set of states $u \in S$ that have a weak transition $u \Rightarrow_c \mu$ where $\text{supp}(\mu) \subseteq Sat(\phi)$.

Distribution formulas are interpreted over distributions $\mu \in \text{Dist}(S)$. The propositional logic fragment has the standard semantics. The meaning of $[\phi]_{\bowtie p}$ is given by:

$$\mu \models [\phi]_{\bowtie p} \quad \text{iff} \quad \mu(Sat(\phi)) \bowtie p$$

Remark 6 (Uniform treatment of modalities for visible and silent actions) Traditionally, Hennessy/Milner-like logics differentiate between modalities dealing with visible actions and the modality for the silent step, while our logics do not. For instance, if considering non-probabilistic branching bisimulation [23,26], a modality $\langle\psi_1\langle\epsilon\rangle\psi_2\rangle$ might be expected with the following semantics:

$$s \models \langle\psi_1\langle\epsilon\rangle\psi_2\rangle \quad \text{iff} \quad \begin{cases} \exists\mu_1, \mu_2 \in \text{Dist}(S) \text{ such that:} \\ \quad 1.\ s \Rightarrow_c \mu_1 \\ \quad 2.\ \mu_1 \xrightarrow{\tau}_c \mu_2 \text{ or } \mu_1 = \mu_2 \\ \quad 3.\ \mu_1 \models \psi_1, \text{ and } \mu_2 \models \psi_2 \end{cases}$$

However, in our setting we consider $s \models \langle\psi_1\langle\epsilon\rangle\psi_2\rangle$ iff $s \models \langle\psi_1\langle\tau\rangle\psi_2\rangle$. This is due to the fact that $\mu_1 = \mu_2$ is redundant since $\mu_1 \xrightarrow{\tau}_c \mu_1$ is always a compound transition (taking the skip probability to be 1).

State-equivalences induced by the logics Let $R(\mathcal{L}^{state}_{x,\xi})$ be the relation on states induced by the logic $\mathcal{L}^{state}_{x,\xi}$ and $R(\mathcal{L}^{dist}_{x,\xi})$ be the relation on distributions induced by the logic $\mathcal{L}^{dist}_{x,\xi}$. That is

$$(s,t) \in R(\mathcal{L}^{state}_{x,\xi}) \quad \text{iff} \quad \text{for all } \phi \in \mathcal{L}^{state}_{x,\xi} : s \models \phi \;\Leftrightarrow\; t \models \phi$$
$$(\mu,\nu) \in R(\mathcal{L}^{dist}_{x,\xi}) \quad \text{iff} \quad \text{for all } \psi \in \mathcal{L}^{dist}_{x,\xi} : \mu \models \psi \Leftrightarrow \nu \models \psi$$

Notice that both $R(\mathcal{L}^{state}_{x,\xi})$ and $R(\mathcal{L}^{dist}_{x,\xi})$ obviously are equivalence relations.

For every equivalence class $C \in S/R(\mathcal{L}^{state}_{x,\xi})$ there exists a formula of $\mathcal{L}^{state}_{x,\xi}$ that distinguishes all states in C from the rest. Similarly for $\mathcal{L}^{dist}_{x,\xi}$. This is stated in the following lemma.

Lemma 15 (Characteristic formulas) *Take any $x \in \{b, \eta, d, w\}$ and $\xi \in \{\nabla, \Delta, s, 0, \lambda, \varnothing\}$. Then:*

(a) *For each $C \in S/R(\mathcal{L}^{state}_{x,\xi})$ there exists $\hat{\phi}_C \in \mathcal{L}^{state}_{x,\xi}$ with $Sat\big(\hat{\phi}_C\big) = C$.*

(b) *$(\mu,\nu) \in R(\mathcal{L}^{dist}_{x,\xi})$ if and only if $\mu \equiv_{R(\mathcal{L}^{state}_{x,\xi})} \nu$.*

(c) *For each $D \in \text{Dist}(S)/R(\mathcal{L}^{dist}_{x,\xi})$ there exists $\hat{\psi}_D \in \mathcal{L}^{dist}_{x,\xi}$ with $Sat\big(\hat{\psi}_D\big) = D$.*

Proof We first observe that for all equivalence classes $C, C' \in S/R(\mathcal{L}^{state}_{x,\xi})$ with $C \neq C'$, there must be a state formula $\phi_{C,C'} \in \mathcal{L}^{state}_{x,\xi}$ that distinguishes the states in C from the states in C'. Because $\mathcal{L}^{state}_{x,\xi}$ contains negation, we can assume that $s \models \phi_{C,C'}$ for all $s \in C$ and $s' \not\models \phi_{C,C'}$ for all $s' \in C'$. For $C \in S/R(\mathcal{L}^{state}_{x,\xi})$, define:

$$\hat{\phi}_C \overset{\text{def}}{=} \bigwedge_{\substack{C' \in S/R(\mathcal{L}^{state}_{x,\xi}) \\ C \neq C'}} \phi_{C,C'}$$

With S also $S/R(\mathcal{L}^{state}_{x,\xi})$ is finite. So, $\hat{\phi}_C \in \mathcal{L}^{state}_{x,\xi}$. Moreover, $s \models \hat{\phi}_C$ iff $s \in C$. That is, $Sat\big(\hat{\phi}_C\big) = C$ which proves statement (a).

For the left to right implication of statement (b) suppose by contradiction that there exists $C \in S/R(\mathcal{L}^{state}_{x,\xi})$ such that $\mu(C) \neq \nu(C)$. Using (a), $\mu \models [\hat{\phi}_C]_{=\mu(C)}$ but $\nu \not\models [\hat{\phi}_C]_{=\mu(C)}$. For the right to left implication notice that every formula in $\mathcal{L}^{dist}_{x,\xi}$ can be written in CNF by $\bigwedge_i \bigvee_j \ell_{ij}$ where each literal ℓ_{ij} has the form $[\phi_{ij}]_{\bowtie q_{ij}}$ or $\neg[\phi_{ij}]_{\bowtie q_{ij}}$. Therefore, it suffices to prove that $\mu(Sat(\phi)) = \nu(Sat(\phi))$ for all $\phi \in \mathcal{L}^{state}_{x,\xi}$. But this is an immediate consequence of $\mu \equiv_{R(\mathcal{L}^{state}_{x,\xi})} \nu$ after observing that $Sat(\phi)$ is a union of $R(\mathcal{L}^{state}_{x,\xi})$ equivalence classes.

To show statement (c) let $D \in \mathrm{Dist}(S)/R(\mathcal{L}^{dist}_{x,\xi})$. Because of (b), for each $C \in S/R(\mathcal{L}^{state}_{x,\xi})$, there exists $q_C \in [0,1]$ such that $\mu(C) = q_C$ for all distributions $\mu \in D$. Hence, $\mu \models [\hat{\phi}_C]_{=q_C}$ for $\mu \in D$.

We now consider the distribution formula

$$\hat{\psi}_D \overset{\text{def}}{=} \bigwedge_{C \in S/R} [\hat{\phi}_C]_{=q_C}$$

Then, $\mu \models \hat{\psi}_D$ iff $\mu(C) = q_C$ for all $C \in S/R(\mathcal{L}^{state}_{x,\xi})$ iff $\mu \in D$. Hence, $Sat\big(\hat{\psi}_D\big) = D$. □

Remark 7 (Finite vs. countable state spaces) Lemma 15 relies on our default assumption that the state space of the given MDP is finite. This implies that the index of any equivalence relation of S is finite and the well-definedness of the formulas $\hat{\phi}_C$ as the conjunctions of the formulas $\phi_{C,C'}$ in the proof of statement (a) in Lemma 15. This argument is no longer adequate for countable state spaces, in which case we would need countable (rather than binary) conjunctions in the logic to ensure the existence of characteristic state formulas.

Remark 8 (Real vs. rational thresholds in distribution formulas) The existence of characteristic formulas for distribution formulas heavily makes use of real threshold values in distribution formulas of the type $[\phi]_{=q}$. (Recall that the proof of statement (c) in Lemma 15 uses the formulas $[\hat{\phi}_C]_{=q_C}$ where $q_C = \mu(C)$). When restricting to rational threshold values in the probability operator would ensure the recursive enumerability of formulas, but we would loose the existence of characteristic formulas for distributions (i.e., statement (c) in Lemma 15 would not hold anymore). Nevertheless, the existence of characteristic state formulas (statement (a) in Lemma 15) as well as statement (b) in Lemma 15 would still hold. The proof of (a) is not affected by the restriction to rational threshold values. The proof of (b) would need to be rephrased as follows.

Given $C \in S/R(\mathcal{L}^{state}_{x,\xi})$ and distributions μ, ν such that $\mu(C) < \nu(C)$ we can pick a rational value q such that $\mu(C) < q < \nu(C)$. But then $[\hat{\phi}_C]_{>q}$ is a distribution formula that distinguishes μ and ν as we have $\mu \not\models [\hat{\phi}_C]_{>q}$, while $\nu \not\models [\hat{\phi}_C]_{>q}$.

However, notice that if a countable conjunction is introduced, any probability operator with real threshold can be derived using the probability operators with rational thresholds, the countable conjunction and, if necessary, negation. □

7 Logical characterisation of bisimulations

We now present the first main result stating that the logics of Section 6 indeed provide sound and complete characterisations of ξ-respecting x-bisimilarity for any (x,ξ)-combination where ξ is different from $\nabla_{>0}$ and $\Delta_{>0}$. This result (Theorem 1 as well as Lemma 16 and Lemma 17) holds for arbitrary (possibly infinite) MDPs.

Theorem 1 (Logical characterisation) *For all $x \in \{b, \eta, w, d\}$ and all $\xi \in \{\nabla, \Delta, s, 0, \lambda, \varnothing\}$, $\approx_x^\xi$ agrees with $R(\mathcal{L}_{x,\xi}^{state})$.*

The proof of Theorem 1 follows from the following Lemma 16 (which shows that $\approx_x^\xi$ is finer than $R(\mathcal{L}_{x,\xi}^{state})$) and Lemma 17 (which shows that $R(\mathcal{L}_{x,\xi}^{state})$ is finer than $\approx_x^\xi$).

Lemma 16 (Soundness: preservation of logical properties) *For each $x \in \{b, \eta, d, w\}$, each $\xi \in \{\nabla, \Delta, s, 0, \lambda, \varnothing\}$:*

1. *$s \approx_x^\xi t$ implies that $(s, t) \in R(\mathcal{L}_{x,\xi}^{state})$, and*
2. *$\mu \approx_x^\xi \nu$ implies that $(\mu, \nu) \in R(\mathcal{L}_{x,\xi}^{dist})$.*[4]

Proof Given a stutter-closed ξ-respecting x-bisimulation R, we prove that for all $s, t \in S$ and $\mu, \nu \in \text{Dist}(S)$ and all formulas $\phi \in \mathcal{L}_{x,\xi}^{state}$ and $\psi \in \mathcal{L}_{x,\xi}^{dist}$:

1. if $(s, t) \in R$ and $s \models \phi$, then $t \models \phi$, and
2. if $\mu \equiv_R \nu$ and $\mu \models \psi$, then $\mu \models \psi$,

from which the claim follows. The proof is by structural induction on the syntactic structure of formulas.

So, we proceed by case analysis. Cases $\top$, $\neg\phi$, $\bigwedge_i \phi_i \in \mathcal{L}_{b,\xi}^{state}$ and $\neg\psi$, $\bigwedge_i \psi_i \in \mathcal{L}_{b,\xi}^{dist}$ are straightforward and we omit them here.

Case $\langle \psi_1 \langle \alpha \rangle \psi_2 \rangle$: Suppose $(s, t) \in R$, with R being a ξ-respecting branching bisimulation, and $s \models \langle \psi_1 \langle \alpha \rangle \psi_2 \rangle$. Then, there exist distributions μ_1 and μ_2 such that

$$s \Rightarrow_c \mu_1 \xrightarrow{\alpha}_c \mu_2, \mu_1 \models \psi_1 \text{ and } \mu_2 \models \psi_2.$$

As a consequence of Lemmas 2 and 3. there are distributions ν_1, ν_1', and ν_2 such that

$$t \Rightarrow_c \nu_1 \Rightarrow_c \nu_1' \xrightarrow{\alpha}_c \nu_2$$

where $\mu_1 \equiv_R \nu_1$, $\mu_1 \equiv_R \nu_1'$, and $\mu_2 \equiv_R \nu_2$. Hence $t \Rightarrow_c \nu_1' \xrightarrow{\alpha}_c \nu_2$, and $\nu_1' \models \psi_1$ and $\nu_2 \models \psi_2$ by induction hypothesis. Therefore, $t \models \langle \psi_1 \langle \alpha \rangle \psi_2 \rangle$.

Case $\langle \psi_1 \langle \alpha \rangle\!\rangle \psi_2$: Suppose $(s, t) \in R$, with R being a ξ-respecting η-bisimuation, and $s \models \langle \psi_1 \langle \alpha \rangle\!\rangle \psi_2$. Then, there exist distributions μ_1, μ_2', and μ_2 such that

$$s \Rightarrow_c \mu_1 \xrightarrow{\alpha}_c \mu_2' \Rightarrow_c \mu_2, \quad \mu_1 \models \psi_1, \text{ and } \mu_2 \models \psi_2.$$

By Lemmas 2 and 3, there are distributions ν_1, ν_1', ν_2' and ν_2 such that

$$t \Rightarrow_c \nu_1 \Rightarrow_c \nu_1' \xrightarrow{\alpha}_c \Rightarrow_c \nu_2' \Rightarrow_c \nu_2$$

where $\mu_1 \equiv_R \nu_1$, $\mu_1 \equiv_R \nu_1'$, and $\mu_2 \equiv_R \nu_2$. Hence, $t \Rightarrow_c \nu_1' \xrightarrow{\alpha}_c \Rightarrow_c \nu_2$, and by induction $\nu_1' \models \psi_1$ and $\nu_2 \models \psi_2$. Therefore, $t \models \langle \psi_1 \langle \alpha \rangle\!\rangle \psi_2$.

Case $\langle\!\langle \alpha \rangle \psi \rangle$: Suppose $(s, t) \in R$, with R being a ξ-respecting delay bisimuation, and $s \models \langle\!\langle \alpha \rangle \psi \rangle$. Then exist distributions μ' and μ with $s \Rightarrow_c \mu' \xrightarrow{\alpha}_c \mu$ and $\mu \models \psi$. As a consequence of Lemmas 2 and 3, there are distributions ν' and ν such that

$$t \Rightarrow_c \nu' \Rightarrow_c \xrightarrow{\alpha}_c \nu, \quad \mu' \equiv_R \nu' \text{ and } \mu \equiv_R \nu.$$

Hence, $t \Rightarrow_c \xrightarrow{\alpha}_c \nu$, and $\nu \models \psi$ by induction hypothesis. Thus, $t \models \langle\!\langle \alpha \rangle \psi \rangle$.

[4] Recall that $\mu \approx_x^\xi \nu$ stands for $\mu \equiv_{\approx_x^\xi} \nu$.

Case $\langle\!\langle\alpha\rangle\!\rangle\psi$: Suppose $(s,t) \in R$, with R being a ξ-respecting weak bisimuation, and $s \models \langle\!\langle\alpha\rangle\!\rangle\psi$. Then, there are distributions μ', μ'' and μ with

$$s \Rightarrow_c \mu' \xrightarrow{\alpha}_c \mu'' \Rightarrow_c \mu \text{ and } \mu \models \psi.$$

As a consequence of Lemmas 2 and 3, there are distributions ν', ν'', and ν such that

$$t \Rightarrow_c \nu' \Rightarrow_c \xrightarrow{\alpha}_c \Rightarrow_c \nu'' \Rightarrow_c \nu$$

where $\mu' \equiv_R \nu'$, $\mu'' \equiv_R \nu''$, and $\mu \equiv_R \nu$. Hence, $t \Rightarrow_c \xrightarrow{\alpha}_c \Rightarrow_c \nu$, and $\nu \models \psi$ by induction hypothesis. Therefore, $t \models \langle\!\langle\alpha\rangle\!\rangle\psi$.

Case $[\phi]_{\bowtie q}$: Suppose $\mu \equiv_R \nu$ and $\mu \models [\phi]_{\bowtie q}$. Then $\mu(Sat(\phi)) \bowtie q$. Since $(s,t) \in R$, by induction hypothesis, $s \models \phi$ implies $t \models \phi$. As a consequence, the $Sat(\phi)$ can be written as union of equivalence classes $C_i \subseteq S/R$, $i \in I$. Hence:

$$\nu(Sat(\phi)) = \nu(\bigcup_{i \in I} C_i) = \sum_{i \in I} \nu(C_i)$$

$$\overset{(*)}{=} \sum_{i \in I} \mu(C_i) = mu(\bigcup_{i \in I} C_i)$$

$$= \mu(Sat(\phi)) \bowtie q$$

where equality $(*)$ holds because $\mu \equiv_R \nu$. Therefore, $\nu \models [\phi]_{\bowtie q}$.

Case $\xi \in \{\Delta, \mathsf{s}, 0, \lambda\}$: Suppose $(s,t) \in R$ and $s \models \xi$. Then $s \Rightarrow_c \mu$ and $\xi(\mu)$, that is, $\mu(\{s' \mid \xi(s')\}) = 1$. By Lemma 2, $t \Rightarrow_c \nu$ and $\mu \equiv_R \nu$ for some distribution ν. Hence, for all states $t' \in \text{supp}(\nu)$ there is some state $s' \in \text{supp}(\mu)$ such that $(s',t') \in R$ and $\xi(s')$. Therefore, for all states $u \in \text{supp}(\nu)$, there is a distribution θ_u with $u \Rightarrow_c \theta_u$ and $\xi(\theta_u)$. With

$$\theta \overset{\text{def}}{=} \sum_{u \in \text{supp}(\nu)} \nu(u) \cdot \theta_u$$

we get $\nu \Rightarrow_c \theta$ and, as a consequence, $t \Rightarrow_c \theta$. Notice also that $\xi(\theta)$ as $\text{supp}(\theta)$ is the union of the supports of the distributions θ_u for $u \in \text{supp}(\nu)$. Frome here we conclude that $t \models \xi$.

Case $\nabla\phi$*for* $\xi = \nabla$ *and* $x \in \{b, \eta\}$: The task is to prove that $\approx_x^\nabla$-equivalent states agree on formulas of the form $\nabla\phi$. Suppose R is a stutter-closed ∇-respecting x-bisimulation.

By induction hypothesis, $Sat(\phi)$ is the union of R-equivalence classes. As before, state s is said to be R-divergent if $\nabla^R(s)$ holds. An infinite path π is called R-divergent if it consists of τ-actions and all states in π belong to the same R-equivalence class, in which case we write $\pi \models div_R$. Pick a scheduler σ_R such that $\Pr_t^{\sigma_R}(div_R) = 1$ for all R-divergent states t. Note that such a scheduler exists, even a deterministic memoryless scheduler. See Corollary 6.

Suppose now that s is a state with $s \models \nabla\phi$. In particular, $s \models \phi$. By induction hypothesis, all states $t \in [s]_R$ satisfy ϕ.

In the simple case where s is R-divergent, then so are all states in the R-equivalence class of s (recall that for ∇-respecting relations R, $\nabla^R(s)$ implies $\nabla^R(t)$ for all $t \in [s]_R$). Thus, the above scheduler σ_R witnesses that all states in $[s]_R$ satisfy $\nabla\phi$.

To treat the general case where a given state s satisfying $\nabla\phi$ might or might not be R-divergent, we regard the set $W = Sat(\nabla\phi)$ and show that W is R-closed, i.e., $s \in W$ and $(s,t) \in R$ implies $t \in W$. Let

$$[W]_R = \left\{ t \in S \mid \exists s \in W \text{ s.t. } (s,t) \in R \right\}$$

be the R-closure of W. As all states $s \in W$ are ϕ-states, the induction hypothesis applied to ϕ yields that $[W]_R \subseteq \mathit{Sat}(\phi)$. To prove $[W]_R \subseteq W$, it suffices to show that each state $t \in [W]_R$ has a transition $t \xrightarrow{\tau} \rho$ in $\mathcal{M}$ where $\mathrm{supp}(\rho) \subseteq [W]_R$.

Pick a state $t \in [W]_R$. By definition of $[W]_R$, there is a state $s \in W$ with $(s, t) \in R$. If s is ∇-diverging then so is t, and the existence of such a τ-transition from t is obvious. Suppose now that s is not ∇-diverging. The scheduler witnessing $s \models \nabla\phi$ yields the existence of distributions μ, μ' and a state $s' \in [s]_R$ such that $s \Longrightarrow_c^R \mu$, $s' \in \mathrm{supp}(\mu)$, $s' \xrightarrow{\tau} \mu'$ $\mathrm{supp}(\mu') \subseteq [W]_R$ and $\mathrm{supp}(\mu')\backslash[s]_R$ is nonempty. The latter implies $\mu' \not\equiv_R \delta_s$. As R is an equivalence, states s' and t are R-equivalent. Hence, t can mimick the transition $s' \xrightarrow{\tau} \mu'$ by concatenating an R-stutter step $t \Longrightarrow_c^R \nu$ with a transition $\nu \xrightarrow{\tau}_c \nu'$ where $\mu' \equiv_R \nu'$ (see Corollary 5). The latter implies $\mathrm{supp}(\nu') \subseteq [W]_R$. But then t has at least one transition $t \xrightarrow{\tau} \rho$ where $\mathrm{supp}(\rho) \subseteq [W]_R$, be it a R-stutter transition or, if $\nu = \delta_t$, the τ-transitions used to generate the compound transition $t \xrightarrow{\tau}_c \nu'$.

Case $\nabla_\epsilon\psi$ *for* $\xi = \nabla$ *and* $x \in \{b, \eta, w, d\}$: As in the previous case we pick a stutter-closed ∇-respecting x-bisimulation R and show that all R-equivalent states have the same truth value for the formula $\nabla_\epsilon\phi$.

Let $U = U_\phi$ denote the set of all states $u \in S$ such that u has a weak transition $u \Rightarrow_c \mu_u$ where $\mathrm{supp}(\mu_u) \subseteq \mathit{Sat}(\phi)$. By induction hypothesis and Lemma 2, $\mathit{Sat}(\phi)$ and U are R-closed.

Let now W denote the largest subset of U such that all states $s \in W$ have a transition $s \xrightarrow{\tau} \nu$ where $\mathrm{supp}(\nu) \subseteq W$. Then, $\mathit{Sat}(\nabla_\epsilon\phi) = W$. The task is to show that W is R-closed.

We define $[W]_R$ as the set of states that are R-equivalent to some state in W. So, the goal is to show that $[W]_R \subseteq W$. For this, it suffices to prove that each state $t \in [W]_R$ has a transition $t \xrightarrow{\tau} \nu$ with $\mathrm{supp}(\nu) \subseteq [W]_R$.

Let $t \in [W]_R$. Then, there is some state $s \in W$ with $(s, t) \in R$.

Again, the case where s is R-divergent is simple. In this case, t is R-divergent as well. Moreover, $t \in U$ (as $s \in U$ and U is R-closed as stated above). Thus, each scheduler witnessing the R-divergence of t is also a witness for $t \models \nabla_\epsilon\phi$. Hence, $t \in W$. In particular, there is a transition $t \xrightarrow{\tau} \nu$ with $\mathrm{supp}(\nu) \subseteq [W]_R$.

Suppose now that s is not R-divergent. Let σ be a scheduler witnessing the satisfaction of $\nabla_\epsilon\phi$ in state s. As σ schedules only τ-transitions and s is not R-divergent, σ induces a weak transition of the form

$$s \Longrightarrow_c^R \theta \xrightarrow{\tau}_c \rho$$

where all involved states belong to W (more precisely, if $\mathcal{T}$ denotes the corresponding τ-tree then $\mathit{state}(v) \in W$ for all nodes v in $\mathcal{T}$) and $\mathrm{supp}(\rho)\backslash[s]_R$ is nonempty.

We may even assume that there is a state $u_0 \in \mathrm{supp}(\theta)$ and a transition $u_0 \xrightarrow{\tau} \rho'$ such that $\rho(u) = \theta(u) + \theta(u_0) \cdot \rho'(u)$ for all states $u \in S\backslash\{u_0\}$ and $\rho(u_0) = \rho'(u_0)$.

As $\theta \equiv_R \delta_s$, state u_0 is R-equivalent to s and t. Thus, there are distributions θ', ρ' with

$$t \Longrightarrow_c^R \theta' \xrightarrow{\tau}_c \rho'$$

where $\rho \equiv_R \rho'$. As $\mathrm{supp}(\rho) \subseteq W \subseteq [W]_R$ and $[W]_R$ is R-closed we obtain $\mathrm{supp}(\rho') \subseteq [W]_R$. This yields the existence of a τ-transition $t \xrightarrow{\tau} \nu$ with $\mathrm{supp}(\nu) \subseteq [W]_R$.

$\square$

Next, we show that $R(\mathcal{L}^{state}_{x,\xi})$ is finer than $\approx^{\xi}_{x}$. For this, we show that $R(\mathcal{L}^{state}_{x,\xi})$ is a ξ-respecting x-bisimulation.

Lemma 17 (Completeness) *Let $x \in \{b, \eta, d, w\}$ and $\xi \in \{\nabla, \Delta, \mathsf{s}, 0, \lambda, \varnothing\}$. Then, the equivalence $R(\mathcal{L}^{state}_{x,\xi})$ is a ξ-respecting x-bisimulation.*

Proof Let $\mathcal{R}_{x,\xi} = R(\mathcal{L}^{state}_{x,\xi})$. Clearly, $\mathcal{R}_{x,\xi}$ is an equivalence relation.

Claim 1. $\mathcal{R}_{x,\xi}$ is an x-bisimulation.

Proof of claim 1. We address here only the case $x = b$. The argument for $x \in \{\eta, w, d\}$ is analogous and omitted here. Suppose that $(s, t) \in \mathcal{R}_{b,\xi}$ and consider first the case $s \xrightarrow{\alpha} \mu$. Let $\hat{\psi}_{[\delta_s]}$ and $\hat{\psi}_{[\mu]}$ be the characteristic formulas of the $R(\mathcal{L}^{dist}_{b,\xi})$-equivalence classes of δ_s and μ (see part (c) of Lemma 15).

Then $s \models \langle \hat{\psi}_{[\delta_s]} \langle\alpha\rangle \hat{\psi}_{[\mu]} \rangle$. Because $(s, t) \in \mathcal{R}_{b,\xi}$, also $t \models \langle \hat{\psi}_{[\delta_s]} \langle\alpha\rangle \hat{\psi}_{[\mu]} \rangle$ and hence $t \Rightarrow_c \nu \xrightarrow{\alpha}_c \nu'$ with $\nu \models \hat{\psi}_{[\delta_s]}$ and $\nu' \models \hat{\psi}_{[\mu]}$. Therefore, using Lemma 15, $\delta_s \equiv_{\mathcal{R}_{b,\xi}} \nu$ and $\mu \equiv_{\mathcal{R}_{b,\xi}} \nu'$, which proves this case.

Claim 2. $\mathcal{R}_{x,\xi}$ is stutter-closed.

Proof of claim 2. Suppose $s \lhd_{\mathcal{R}_{x,\xi}} t$. By Lemma 4 there are distributions μ, ν such that $s \Rightarrow_c \nu$ and $t \Rightarrow_c \mu$ with $\mu \equiv_{\mathcal{R}_{x,\xi}} \delta_s$ and $\nu \equiv_{\mathcal{R}_{x,\xi}} \delta_t$. By inspection on the semantics of $\mathcal{L}^{state}_{x,\xi}$ we obtain the $\mathcal{R}_{x,\xi}$-equivalence of s and t.

Claim 3. $\mathcal{R}_{x,\xi}$ is ξ-respecting.

Proof of claim 3. Let us first consider the cases $\xi \in \{\Delta, \mathsf{s}, 0, \lambda\}$. Suppose s is a state with $\xi(s)$. Then, $s \models \xi$ and hence also $t \models \xi$ for each state t in the $R(\mathcal{L}^{state}_{x,\xi})$-equivalence class of s. From here $t \Rightarrow_c \nu$ and $\xi(\nu)$ for some ν, which proves this case.

We now address the case $\xi = \nabla$. We start with the case $x = b$. The argument for $x = \eta$ is analogous. Suppose s is a state such that $\nabla^{\mathcal{R}_{b,\nabla}}(s)$. Let C denote the $\mathcal{R}_{b,\nabla}$-equivalence class of s and $\hat{\phi}_C$ its characteristic formula (see part (a) of Lemma 15). Then, there is a scheduler σ such that $\Pr^{\sigma}_{s}\left((C \times \{\tau\})^{\omega}\right) = 1$. As $C = Sat(\hat{\phi}_C)$ we obtain:

$$\Pr^{\sigma}_{s}\left(\left(Sat(\hat{\phi}_C) \times \{\tau\}\right)^{\omega}\right) = 1$$

and hence $s \models \nabla\hat{\phi}_C$. If $(s, t) \in \mathcal{R}_{b,\nabla}$ then also $t \models \nabla\hat{\phi}_C$. Since $t \in C$ we conclude that $\nabla^{\mathcal{R}_{b,\nabla}}(t)$.

It remains to consider the case $\xi = \nabla$ and $x \in \{w, d\}$. Suppose s is a state with $\nabla^{\mathcal{R}_{x,\nabla}}(s)$ and let t be a state that is $\mathcal{R}_{x,\nabla}$-equivalent to s. As before, let C denote the $\mathcal{R}_{x,\nabla}$-equivalence class of s (and t) and $\hat{\phi}_C$ its characteristic formula (see part (a) of Lemma 15). Obviously, $\nabla^{\mathcal{R}_{x,\nabla}}(s)$ implies $s \models \nabla_{\epsilon}\hat{\phi}_C$. Therefore, $t \models \nabla_{\epsilon}\hat{\phi}_C$. So, there is a scheduler σ such that

$$\Pr^{\sigma}_{t}\left(\left(T \times \{\tau\}\right)^{\omega}\right) = 1$$

where T consists of all states $u \in S$ such that u has a weak transition $u \Rightarrow_c \mu$ where $\mathrm{supp}(\mu) \subseteq Sat(\hat{\phi}_C) = C$. But then t is $\mathcal{R}_{x,\nabla}$-divergent by Lemma 8, i.e., $\nabla^{\mathcal{R}_{x,\nabla}}(t)$ holds. □

Remark 9 (Coarsest bisimulation) The logical characterisation in Theorem 1 (showing that $\approx^{\xi}_{x}$ equals $R(\mathcal{L}^{state}_{x,\xi})$) together with Lemma 17 (showing that $R(\mathcal{L}^{state}_{x,\xi})$ is a ξ-respecting x-bisimulation) yields an alternative proof for the statement that $\approx^{\xi}_{x}$ is the coarsest ξ-respecting x-bisimulation as stated in Lemma 7.

Logical characterisations of $\Delta_{>0}$- and $\nabla_{>0}$-respecting bisimulations

In Remarks 2 and 3, we introduced the predicates $\nabla_{>0}$ and $\Delta_{>0}$, and showed that $\nabla_{>0}$-respecting x-bisimilarity $\approx_x^{\nabla_{>0}}$ is strictly coarser than ∇-respecting x-bisimilarity $\approx_x^{\nabla}$, together with the analogous result for $\Delta_{>0}$ and Δ. This raises the question for logical characterisations of ξ-respecting x-bisimilarity $\approx_x^{\xi}$ for $\xi \in \{\nabla_{>0}, \Delta_{>0}\}$. While such a characterisation is open for the two combinations $x \in \{w, d\}$ and $\xi = \nabla_{>0}$, logical characterisations for $\approx_b^{\nabla_{>0}}$, $\approx_\eta^{\nabla_{>0}}$ in finite MDPs and $\approx_x^{\Delta_{>0}}$ where $x \in \{b, \eta, w, d\}$ in arbitrary (possibly infinite MDPs) are obtained as follows.

$\Delta_{>0}$-respecting x-bisimilarity. For $x \in \{b, \eta, w, d\}$, let $\mathcal{L}^{state}_{x,\Delta_{>0}}$ and $\mathcal{L}^{dist}_{x,\Delta_{>0}}$ be defined as $\mathcal{L}^{state}_{x,\varnothing}$ and $\mathcal{L}^{dist}_{x,\varnothing}$ by adding the atomic state formula $\Delta_{>0}$ with the semantics $s \models \Delta_{>0}$ iff $\Delta_{>0}(s)$. Using that $R(\mathcal{L}^{state}_{x,\varnothing})$ agrees with $\approx_x^{\varnothing}$ (Theorem 1), we obtain that the equivalence $R(\mathcal{L}^{state}_{x,\Delta_{>0}})$ that identifies all states satisfying the same $\mathcal{L}^{state}_{x,\Delta_{>0}}$ formulas agrees with $\approx_x^{\Delta_{>0}}$.

$\nabla_{>0}$-respecting branching and η-bisimilarity. For $x \in \{b, \eta\}$ and explicit divergence, we deal with the logic $\mathcal{L}^{state}_{x,\nabla_{>0}}$ and $\mathcal{L}^{dist}_{x,\nabla_{>0}}$ that agree with $\mathcal{L}^{state}_{x,\nabla}$ and $\mathcal{L}^{dist}_{x,\nabla}$ except that we replace the formulas $\nabla\phi$ with $\nabla_{>0}\phi$. The semantics of this new operator is given by:

$$s \models \nabla_{>0}\phi \;\text{ iff }\; \exists\sigma \text{ s.t. } \Pr_s^\sigma\big(\, (Sat(\phi) \times \{\tau\})^\omega \,\big) \;>\; 0$$

As before, we write $R(\mathcal{L}^{state}_{x,\nabla_{>0}})$ and $R(\mathcal{L}^{dist}_{x,\nabla_{>0}})$ for the induced equivalence on states and distributions, respectively.

Lemma 18 *For each $x \in \{b, \eta\}$, the relations $R(\mathcal{L}^{state}_{x,\nabla_{>0}})$ and $\approx_x^{\nabla_{>0}}$ coincide in finite MDPs.*

Proof We show that $\approx_x^{\nabla_{>0}}$ is finer than $R(\mathcal{L}^{state}_{x,\nabla_{>0}})$, and vice versa.

Claim 1: $\approx_x^{\nabla_{>0}}$ is finer than $R(\mathcal{L}^{state}_{x,\nabla_{>0}})$.

Proof of claim 1. To prove that $\approx_x^{\nabla_{>0}}$ preserves the truth value of all $\mathcal{L}^{state}_{x,\nabla_{>0}}$ formulas we extend the proof of Lemma 16 for the case $\xi = \nabla_{>0}$. For this we pick a stutter-closed $\nabla_{>0}$-respecting x-bisimulation R and suppose that ϕ is a state formula where $Sat(\phi)$ is the union of R-equivalence classes (induction hypothesis). In what follows, we say state s is positively R-divergent if $\sup \mathbb{DP}_\nabla(s, R) > 0$. Thus, being $\nabla_{>0}$-respecting means that either all states in an R-equivalence class are positively R-divergent or none of them.

Suppose s, t are states such that $(s, t) \in R$ and $s \models \nabla_{>0}\phi$. We aim to prove that $t \models \nabla_{>0}\phi$. For this, we pick a scheduler σ such that $\Pr_s^\sigma\big((Sat(\phi) \times \{\tau\})^\omega\big)$ is positive. Then, $s \models \phi$ and therefore $C \subseteq Sat(\phi)$ where $C = [s]_R = [t]_R$.

- If $\Pr_s^\sigma\big(\, (C \times \{\tau\})^\omega \,\big) \;>\; 0$ then s is positively R-divergent, and hence, so is t. As $C \subseteq Sat(\phi)$, we get $t \models \nabla_{>0}\phi$.
- Suppose now that $\Pr_s^\sigma\big((C \times \{\tau\})^\omega\big) = 0$. Then, we pick a finite σ-path $\pi = s_0\,\tau\,s_1\,\tau \ldots \tau\,s_n$ from $s = s_0$ to a state s_n that belongs to an end component $\mathcal{E}$ consisting of states in $Sat(\phi)$ and τ-transitions and that enjoys the property that $s_i \models \phi$ for $i = 0, 1, \ldots, n$. But then state s_n is R-divergent, in particular positively R-divergent. Hence, $(s, s_n) \notin R$. Let $i_1 < i_2 < \cdots < i_\ell$ be integers in $\{0, \ldots, n\}$ such that $i_1 = 0$ and $(s_j, s_{j+1}) \in R$ for all $j \in \mathbb{N}$ with $i_h \leqslant j < i_{h+1}$ for $h = 1, \ldots, \ell$ where $i_{\ell+1} = n+1$. Let C_h denote the R-equivalence class of states $s_{i_h}, s_{i_h+1}, \ldots, s_{i_{h+1}-1}$.
  As s and t are R-equivalent, there exists a finite path $\pi' = t_0\,\tau\,t_1\,\tau \ldots \tau\,t_m$ from $t = t_0$ that belongs to $(C_1 \times \{\tau\})^+(C_2 \times \{\tau\})^+ \ldots (C_\ell \times \{\tau\})^+$. As all states in π are ϕ-states,

we have $C_h \subseteq Sat(\phi)$ for $h = 1, \dots, \ell$. In particular, all states in π' are ϕ-states. Moreover, as state s_n is positively R-divergent and $C_\ell = [s_n]_R = [t_m]_R$, state t_m is positively R-divergent too. (Here, we use the assumption that R is $\nabla_{>0}$-respecting.) Let σ_0 be a scheduler witnessing $u \models \nabla_{>0}\phi$ for all states $u \in Sat(\nabla_{>0}\phi)$. Now pick a scheduler σ' such that π' is a σ'-path and that behaves as σ_0 after having generated π'. Then, $\Pr_t^{\sigma'}\big((Sat(\phi) \times \{\tau\})^\omega\big) > 0$, and therefore $t \models \nabla_{>0}\phi$.

Claim 2: $R(\mathcal{L}^{state}_{x,\nabla_{>0}})$ is finer than $\approx_x^{\nabla_{>0}}$.

Proof of claim 2. We show that $\mathcal{L}^{state}_{x,\nabla_{>0}}$ and $\mathcal{L}^{dist}_{x,\nabla_{>0}}$ provide a complete characterisation of $\approx_x^{\nabla_{>0}}$. For this, we extend the proof of Lemma 17 for the case $\xi = \nabla_{>0}$.

Let now $\mathcal{R}_{x,\nabla_{>0}}$ denote the equivalence induced by $\mathcal{L}^{state}_{x,\nabla_{>0}}$, i.e., $\mathcal{R}_{x,\nabla_{>0}} = R(\mathcal{L}^{state}_{x,\nabla_{>0}})$. As for the other cases of ξ we get that $\mathcal{R}_{x,\nabla_{>0}}$ is a stutter-closed x-bisimulation (see Claim 1 and 2 in the proof of Lemma 17). The remaining task is to show that $\mathcal{R}_{x,\nabla_{>0}}$ is $\nabla_{>0}$-respecting. For this we first observe that the existence of characteristic formulas as in Lemma 15 also holds for $\xi = \nabla_{>0}$. Let now $(s,t) \in \mathcal{R}_{x,\nabla_{>0}}$ and suppose s is positively $\mathcal{R}_{x,\nabla_{>0}}$-divergent. The goal is to show that t is positively $\mathcal{R}_{x,\nabla_{>0}}$-divergent too.

Pick a characteristic formula $\hat{\phi}_C$ of the $\mathcal{R}_{x,\nabla_{>0}}$-equivalence class C of s and t. That is, $\hat{\phi}_C \in \mathcal{L}^{state}_{x,\nabla_{>0}}$ and $Sat(\hat{\phi}_C) = C$. Then, $s \models \nabla_{>0}\hat{\phi}_C$ and therefore $t \models \nabla_{>0}\hat{\phi}_C$. But then t is positively $\mathcal{R}_{x,\nabla_{>0}}$-divergent. □

8 The spectrum

In this section we discuss the relation between the expressiveness of the logics and between the semantics equivalences.

Since the modalities for actions and divergence are the ones that make the difference in the logics, Lemma 20 will provide encodings of these modalities in terms of operators of the stronger logics. Before, notice that the distribution formula $[\top]_{\geq 0}$ holds in any probability distribution, i.e., $\mu \models [\top]_{\geq 0}$ for any μ. Thus, overloading the $\top$ symbol, we extend the tautology operator to distributions by defining it as $[\top]_{\geq 0}$.

We first state an auxiliary lemma to ease notation.

Lemma 19 *The following equivalences hold:*

$$\langle\!\langle \tau \rangle\!\rangle \psi \;\equiv\; \langle\!\langle \tau \rangle \psi \rangle \;\equiv\; \langle \top \langle \tau \rangle\!\rangle \psi \;\equiv\; \langle \top \langle \tau \rangle \psi \rangle$$
$$\langle\!\langle a \rangle\!\rangle \top \;\equiv\; \langle\!\langle a \rangle \top \rangle \;\equiv\; \langle \top \langle a \rangle\!\rangle \top \;\equiv\; \langle \top \langle a \rangle \top \rangle$$

where the equivalence relation $\equiv$ should be understood as usual: $\phi_1 \equiv \phi_2$ whenever for all state s, $s \models \phi_1$ iff $s \models \phi_2$.

Proof We focus on the proof of $\langle\!\langle \tau \rangle\!\rangle \psi \equiv \langle \top \langle \tau \rangle \psi \rangle$. The other cases follow similarly. Suppose $s \models \langle\!\langle \tau \rangle\!\rangle \psi$. Then there is some μ such that $s \Rightarrow_c \xrightarrow{\tau}_c \Rightarrow_c \mu$ and $\mu \models \psi$. That is, $s \Rightarrow_c \mu$. Then $s \Rightarrow_c \mu \xrightarrow{\tau}_c \mu$ and since $\mu \models \top$, $s \models \langle \top \langle \tau \rangle \psi \rangle$.

Now suppose $s \models \langle \top \langle \tau \rangle \psi \rangle$. Then $s \Rightarrow_c \mu_1 \xrightarrow{\tau}_c \mu_2$, for some μ_1 and μ_2 with $\mu_1 \models \top$, and $\mu_2 \models \psi$. Hence $s \Rightarrow_c \mu_1 \xrightarrow{\tau}_c \mu_2 \Rightarrow_c \mu_2$, and therefore $s \models \langle\!\langle \tau \rangle\!\rangle \psi$. □

Lemma 20 *The following equivalences hold where* ($\Delta_{>0}2$) *relies on the finiteness assumption, while all other statements hold for countable MDPs:*

$(d)\quad \langle\!\langle \alpha \rangle \psi \rangle \;\equiv\; \langle \top \langle \alpha \rangle \psi \rangle \qquad\qquad (\nabla_\epsilon)\quad \nabla_\epsilon \phi \;\equiv\; \nabla \langle\!\langle \tau \rangle\!\rangle [\phi]_{=1}$

$$(w)\quad \langle\!\langle\alpha\rangle\!\rangle\psi \;\equiv\; \langle\top\langle\alpha\rangle\psi \qquad\qquad (\Delta 1)\quad \Delta \;\equiv\; \nabla\top$$
$$(\lambda)\quad \lambda \;\equiv\; \langle\!\langle\tau\rangle\!\rangle[\Delta \vee 0]_{=1} \qquad\qquad (\Delta 2)\quad \Delta \;\equiv\; \nabla_\epsilon\top$$
$$(s)\quad s \;\equiv\; \langle\!\langle\tau\rangle\!\rangle[\neg\Delta_{>0}]_{=1} \qquad\qquad (\Delta_{>0}1)\quad \Delta_{>0} \;\equiv\; \nabla_{>0}\top$$
$$(0)\quad 0 \;\equiv\; \langle\!\langle\tau\rangle\!\rangle[s \wedge \neg\langle\!\langle\tau\rangle\!\rangle[\textstyle\bigvee_{a\neq\tau}\langle\!\langle a\rangle\!\rangle\top]_{>0}]_{=1} \qquad\qquad (\Delta_{>0}2)\quad \Delta_{>0} \;\equiv\; \langle\!\langle\tau\rangle\!\rangle[\Delta]_{>0}$$

Proof The proof of cases (d), (w), (∇_ϵ), $(\Delta 1)$, $(\Delta 2)$, and $(\Delta_{>0}1)$ follow directly by definition. In the following, we prove the remaining cases.

Case $(\Delta_{>0}2)$: The fact that $s \models \langle\!\langle\tau\rangle\!\rangle[\Delta]_{>0}$ implies $s \models \Delta_{>0}$ follows by the definitions. The proof of the other implication follows closely results in Remark 3. So we provide only a sketch of a proof. Notice that for this implication we assume that the MDP is finite. Suppose $s \models \Delta_{>0}$. Then $\Delta_{>0}(s)$ and, as in Lemma 13, there is a memoryless deterministic scheduler $\sigma = \sigma_\Delta^{\max}$ with $\Pr_s^\sigma\big((S\times\{\tau\})^\omega\big) > 0$. Therefore, there is τ-path reachable from s through σ to a BSCC $\mathcal{B}$ built by τ-transitions defined by σ. That is, $s \Rightarrow_c \mu$ and $\mu(\mathcal{B}) > 0$ and for every state t in $\mathcal{B}$, $\Delta(t)$ and hence $t \models \Delta$. As a consequence $s \models \langle\!\langle\tau\rangle\!\rangle[\Delta]_{>0}$.

Case (λ): If $s \models \lambda$, for some μ, $s \Rightarrow_c \mu$ and $\lambda(\mu)$. The later implies that $1 = \mu(\{t \mid \lambda(t)\}) = \mu(\{t \mid \Delta(t) \vee 0(t)\})$. Notice that $\{t \mid \Delta(t) \vee 0(t)\} \subseteq \mathit{Sat}(\Delta \vee 0)$. Thus $\mu(\mathit{Sat}(\Delta \vee 0)) = 1$ and $s \models \langle\!\langle\tau\rangle\!\rangle[\Delta \vee 0]_{=1}$.
Suppose now that $s \models \langle\!\langle\tau\rangle\!\rangle[\Delta \vee 0]_{=1}$. Then, for some μ, $s \Rightarrow_c \mu$ and $\mu(\mathit{Sat}(\Delta \vee 0)) = 1$. The later means that for every $t \in \mathrm{supp}(\mu)$, either $t \models \Delta$ or $t \models 0$. Thus

$$(\exists\mu_t : t \Rightarrow_c \mu_t \text{ and } \mu(\{t \mid \Delta(t)\}) = 1) \text{ or}$$
$$(\exists\mu_t : t \Rightarrow_c \mu_t \text{ and } \mu(\{t \mid 0(t)\}) = 1)\,.$$

Since $\lambda(t)$ iff $\Delta(t)$ or $0(t)$, this is equivalente to

$$\exists\mu_t : t \Rightarrow_c \mu_t \text{ and } \mu(\{t \mid \lambda(t)\}) = 1.$$

Define $\mu' \stackrel{\text{def}}{=} \sum_{t\in\mathrm{supp}(\mu)} \mu(t)\cdot\mu_t$. Then $s \Rightarrow_c \mu \Rightarrow_c \mu'$ and $\lambda(\mu')$. That is, $s \models \lambda$.

Case (s): Suppose $s \models \mathsf{s}$. Then there exists μ s.t. $s \Rightarrow_c \mu$ and $\mathsf{s}(\mu)$. The later implies that for every $t \in \mathrm{supp}(\mu)$, $t \stackrel{\tau}{\not\rightarrow}$, from which $\Pr_t^\sigma\big((S\times\{\tau\})^\omega\big) = 0$ for any possible scheduler σ. Hence $t \models \neg\Delta_{>0}$. Therefore, $\mu \models [\neg\Delta_{>0}]_{=1}$ and, in consequence, $s \models \langle\!\langle\tau\rangle\!\rangle[\neg\Delta_{>0}]_{=1}$.
If $s \models \langle\!\langle\tau\rangle\!\rangle[\neg\Delta_{>0}]_{=1}$, there exists μ s.t. $s \Rightarrow_c \mu$ and $\mu(\mathit{Sat}(\neg\Delta_{>0})) = 1$. Then, for every $t \in \mathrm{supp}(\mu)$, $t \models \neg\Delta_{>0}$, which means that for all scheduler σ, $\Pr_t^\sigma\big((S\times\{\tau\})^\omega\big) = 0$. As a consequence, there must exists μ_t s.t. $t \Rightarrow_c \mu_t$ and $\mu_t(\{t' \mid t' \stackrel{\tau}{\not\rightarrow}\}) = 1$ (otherwise, a possible scheduler could keep choosing τ transitions so that it yields positive probabilities to diverge). That is, $\mathsf{s}(\mu_t)$. Define $\mu' \stackrel{\text{def}}{=} \sum_{t\in\mathrm{supp}(\mu)} \mu(t)\cdot\mu_t$. Then $s \Rightarrow_c \mu \Rightarrow_c \mu'$ and $\mathsf{s}(\mu')$. That is $s \models \mathsf{s}$.

Case (0): If $s \models 0$, there is some μ s.t. $s \Rightarrow_c \mu$ and $\mu(\{t \mid 0(t)\}) = 1$. Thus, if $t \in \mathrm{supp}(\mu)$, $t \stackrel{\alpha}{\not\rightarrow}$ for any $\alpha \in \mathrm{Act}$, which means that $t \models \mathsf{s}$ and $t \models \neg\langle\!\langle\tau\rangle\!\rangle[\bigvee_{a\neq\tau}\langle\!\langle a\rangle\!\rangle\top]_{>0}$. Therefore, $s \models \langle\!\langle\tau\rangle\!\rangle[\mathsf{s} \wedge \neg\langle\!\langle\tau\rangle\!\rangle[\bigvee_{a\neq\tau}\langle\!\langle a\rangle\!\rangle\top]_{>0}]_{=1}$.
Suppose now that $s \models \langle\!\langle\tau\rangle\!\rangle[\mathsf{s} \wedge \neg\langle\!\langle\tau\rangle\!\rangle[\bigvee_{a\neq\tau}\langle\!\langle a\rangle\!\rangle\top]_{>0}]_{=1}$. Then, there exists μ s.t. $s \Rightarrow_c \mu$ and for all $t \in \mathrm{supp}(\mu)$, $t \models \mathsf{s} \wedge \neg\langle\!\langle\tau\rangle\!\rangle[\bigvee_{a\neq\tau}\langle\!\langle a\rangle\!\rangle\top]_{>0}$. That is, for all $t \in \mathrm{supp}(\mu)$,

(i) there exists μ_t such that $t \Rightarrow_c \mu_t$ and $\mu_t(\{t \mid \mathsf{s}(t)\}) = 1$, and
(ii) for all μ_t' such that $t \Rightarrow_c \mu_t'$, $\mu_t'(Sat(\langle\!\langle a \rangle\!\rangle \top)) = 0$ for all $a \neq \tau$.

Because of (ii), $\mu_t(Sat(\langle\!\langle a \rangle\!\rangle \top)) = 0$ for all $a \neq \tau$. Thus, for every $t' \in \mathrm{supp}(\mu_t)$, $t' \Rightarrow_c \stackrel{a}{\not\rightarrow}$. In particular, $t' \stackrel{a}{\not\rightarrow}$ for all $a \neq \tau$. Since in addition $t' \stackrel{\tau}{\not\rightarrow}$ (because of (i)), $0(\mu_t)$. Define $\mu' \stackrel{\text{def}}{=} \sum_{t \in \mathrm{supp}(\mu)} \mu(t) \cdot \mu_t$. Then $s \Rightarrow_c \mu \Rightarrow_c \mu'$ and $0(\mu')$. That is $s \models 0$. □

For the case (0), notice in particular that if there is no $a \neq \tau$ (that is, Act $= \{\tau\}$), $\neg\langle\!\langle \tau \rangle\!\rangle[\bigvee_{a \neq \tau} \langle\!\langle a \rangle\!\rangle \top]_{>0} \equiv \neg\langle\!\langle \tau \rangle\!\rangle[\bot]_{>0} \equiv \top$. Therefore $0 \equiv \langle\!\langle \tau \rangle\!\rangle[\mathsf{s}]_{=1} \equiv \mathsf{s}$, which is indeed as expected since (always under Act $= \{\tau\}$), $0(t)$ iff $\mathsf{s}(t)$ for all state t. Also, contrary to intuition, notice that 0 and $\langle\!\langle \tau \rangle\!\rangle[\mathsf{s} \wedge \bigwedge_{a \neq \tau} \neg\langle\!\langle a \rangle\!\rangle \top]_{=1}$ are note equivalent. Consider the simple MDP with states s, t, u, and transitions $s \xrightarrow{\tau} \frac{1}{2}\delta_t + \frac{1}{2}\delta_u$, $t \xrightarrow{a} \delta_t$, and $u \xrightarrow{a} \delta_u$. Then $s \models \langle\!\langle \tau \rangle\!\rangle[\mathsf{s} \wedge \bigwedge_{a \neq \tau} \neg\langle\!\langle a \rangle\!\rangle \top]_{=1}$ but $s \not\models 0$.

As a corollary of Lemma 20 we have that some logics are more expressive than others. For instance, as a consequence of (d), ($\Delta 1$), ($\Delta_{>0}1$), and (s), $\mathcal{L}^{state}_{b,\nabla}$ is more expressive than $\mathcal{L}^{state}_{d,\mathsf{s}}$.

However, we have failed to make a connection between $\langle \psi_1 \langle \alpha \rangle\!\rangle \psi_2$ and $\langle \psi_1 \langle \alpha \rangle \psi_2 \rangle$, $\langle\!\langle \alpha \rangle\!\rangle \psi$ and $\langle\!\langle \alpha \rangle \psi \rangle$, and $\nabla_{>0}\phi$ and $\nabla\phi$. Therefore we have not established a relation of expressiveness between $\mathcal{L}^{state}_{b,\xi}$ and $\mathcal{L}^{state}_{\eta,\xi}$, $\mathcal{L}^{state}_{d,\xi}$ and $\mathcal{L}^{state}_{w,\xi}$, and $\mathcal{L}^{state}_{x,\nabla}$ and $\mathcal{L}^{state}_{x,\nabla_{>0}}$. Moreover, we suspect that these pairs of logics cannot be expressed in terms of each other (one direction, as expected, is actually known). So, to relate the equivalences induced by these logics, we turn to the operational definitions and prove that a relation R satisfying the stronger requirement, also satisfies the weaker requirements. In the following lemma we also include the cases of the $\nabla_{>0}$-respecting d and w-bisimulations which do not have a logical characterization.

Lemma 21 *Let $\xi \in \{\nabla, \nabla_{>0}, \Delta, \Delta_{>0}, \lambda, \mathsf{s}, 0, \varnothing\}$ and $x \in \{d, w\}$. Then, the following holds.*

(a) *If R is a ξ-respecting b-bisimulation, then it is also a ξ-respecting η-bisimulation.*
(b) *If R is a ξ-respecting d-bisimulation, then it is also a ξ-respecting w-bisimulation.*
(c) *If R is a ∇-respecting x-bisimulation in a finite MDP, then it is also a $\nabla_{>0}$-respecting x-bisimulation.*
(c) *If R is a $\nabla_{>0}$-respecting x-bisimulation in a finite MDP, then it is also a $\Delta_{>0}$-respecting x-bisimulation.*

Proof For (a), since R is ξ-respecting we only have to prove that it is an η-bisimulation. So, let $(s, t) \in R$ and suppose $s \xrightarrow{\alpha} \mu$. Because R is a b-bisimulation, there are of ν and ν' such that $t \Rightarrow_c \nu \xrightarrow{\alpha}_c \nu'$, $\delta_s \equiv_R \nu$, and $\mu \equiv_R \nu'$. Since $\nu' \Rightarrow_c \nu'$, then $t \Rightarrow_c \nu \xrightarrow{\alpha}_c \Rightarrow_c \nu'$ thus proving (a). The case of (b) follows similarly.

Case (c) follows from Lemma 10 and case (d) is stated in Remark 5. □

Theorem 2 *Consider the following (partial) order*

$$\nabla \preceq \Delta \preceq \lambda \preceq \varnothing \qquad \nabla \preceq \nabla_{>0} \preceq \Delta_{>0} \preceq \mathsf{s} \preceq 0 \preceq \varnothing \qquad \Delta \preceq \Delta_{>0}$$
$$b \preceq \eta \preceq w \qquad b \preceq d \preceq w$$

Then $\approx^{\xi_1}_{x_1} \subseteq \approx^{\xi_2}_{x_2}$ whenever $\xi_1 \preceq \xi_2$ and $x_1 \preceq x_2$. Moreover, the inclusion is strict if either $\xi_1 \neq \xi_2$ or $x_1 \neq x_2$.

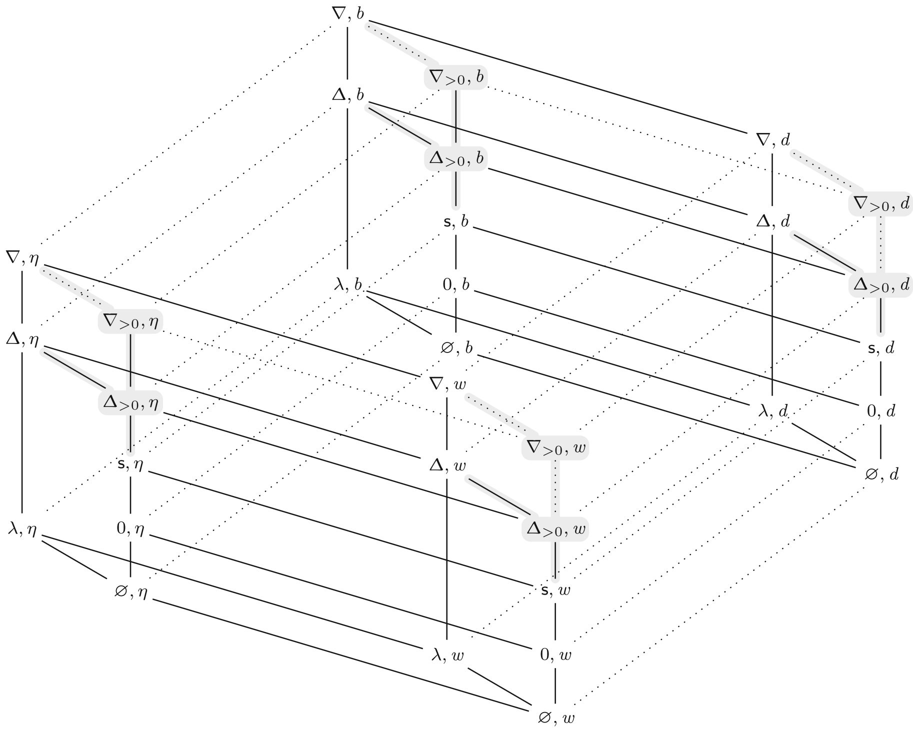

Fig. 6 The probabilistic bisimulation spectrum with silent moves

Proof All the inclusions except when $x_1 \in \{b, d\}$ and $x_2 \in \{\eta, w\}$, when $\xi_1 = \nabla$ and $\xi_2 = \nabla_{>0}$, or when $\xi_i = \nabla_{>0}$ and $x_i \in \{d, w\}$ for any $i \in \{1, 2\}$, are consequence of Lemma 20, Theorem 1, and the results in previous section. The remaining cases are consequences of Lemma 21.

The fact that the inclusions are strict is as follows. The cases where $x_1 \neq x_2$ are well known in the non-probabilistic setting and inherited here (see, e.g., [23,26]). For the cases where $\xi_1, \xi_2 \in \{\nabla, \Delta, \lambda, \mathsf{s}, 0, \varnothing\}$ with $\xi_1 \neq \xi_2$, counterexamples are provided in Fig. 3. The example in Remark 2 shows that $s \approx_{x_1}^{\xi_1} t$ but $s \not\approx x_2\xi_2 t$ whenever $\xi_1 \in \{\nabla_{>0}, \Delta_{>0}\}$ and $\xi_2 \in \{\nabla, \Delta, \lambda\}$. Finally, the two MDPs in the second example in Fig. 3 are ξ_1-respecting x_1-bisimilar but not ξ_2-respecting x_2-bisimilar for $\xi_1 \in \{\mathsf{s}, 0, \varnothing\}$ and $\xi_2 \in \{\nabla_{>0}, \Delta_{>0}\}$, while the fourth example shows the same with $\xi_1 = \Delta_{>0}$ and $\xi_2 = \nabla_{>0}$. □

Figure 6 summarises the results of Theorem 2. Solid lines indicate both strict inclusion of the semantic equivalences and that the logic above is strictly more expressive than the logic below. Dotted lines indicate only inclusion of the semantic equivalences but do not relate the logics. In fact, we conjecture that these logics cannot be expressed one in terms of the other. Grey shades indicate restrictions to finite-state models.

Remark 10 (Alternative characterisation) We could have extended the logics with a simple operator so that the modalities $\langle_\sqcup\langle\alpha\rangle\!\rangle_\sqcup$ and $\langle\!\langle\alpha\rangle\!\rangle_\sqcup$ can be encoded in terms of $\langle_\sqcup\langle\alpha\rangle_\sqcup\rangle$ and $\langle\!\langle\alpha\rangle_\sqcup\rangle$ respectively. To arrive there, consider the new logic operator $\epsilon\psi$ (where ψ is a distribution formula and so is $\epsilon\psi$) whose semantics is defined as follows. If $\mu \in \mathrm{Dist}(S)$ then:

$$\mu \models \epsilon\psi \quad \text{iff} \quad \exists \nu \in \mathrm{Dist}(S) \text{ s.t. } \mu \Rightarrow_c \nu \text{ and } \nu \models \psi$$

Denote by $\mathcal{L}^{state}_{x,\xi}{+}\epsilon$ the extensions of the logic $\mathcal{L}^{state}_{x,\xi}$ that admits this new operator on distribution formulas.

Corollary 1 guarantees that any logic extended with this operator is still sound for its respective bisimulation. Therefore $\approx^{\xi}_{x}$ agrees with $R(\mathcal{L}^{state}_{x,\xi}{+}\epsilon)$ for all $x \in \{b, \eta, d, w\}$ and $\xi \in \{\nabla, \nabla_{>0}, \Delta, \Delta_{>0}, \mathsf{s}, 0, \lambda, \varnothing\}$ except for the combinations of $\xi = \nabla_{>0}$ and $x \in \{w, d\}$, in which case the logics are not defined.

With this new operator, Lemma 20 could be extended by the following two equivalences

$$(\eta)\quad \langle\psi_1\langle\alpha\rangle\!\rangle\psi_2 \;\equiv\; \langle\psi_1\langle\alpha\rangle\epsilon\psi_2\rangle \qquad\qquad (w')\quad \langle\!\langle\alpha\rangle\!\rangle\psi \;\equiv\; \langle\!\langle\alpha\rangle\epsilon\psi\rangle$$

Their respective proof follow easily from the definitions. We briefly show the case of (η). Let $s \models \langle\psi_1\langle\alpha\rangle\!\rangle\psi_2$. Then, there are distributions μ_1, μ_2, and μ_3 such that

$$s \Rightarrow_c \mu_1 \xrightarrow{\alpha}_c \mu_2 \Rightarrow_c \mu_3, \mu_1 \models \psi_1, \text{ and } \mu_3 \models \psi_2.$$

Then, $\mu_2 \models \epsilon\psi_2$ and hence $s \models \langle\psi_1\langle\alpha\rangle\epsilon\psi_2\rangle$.

Conversely, let $s \models \langle\psi_1\langle\alpha\rangle\epsilon\psi_2\rangle$. Then, there are distributions μ_1 and μ_2, such that

$$s \Rightarrow_c \mu_1 \xrightarrow{\alpha}_c \mu_2, \mu_1 \models \psi_1, \text{ and } \mu_2 \models \epsilon\psi_2.$$

Therefore, there is some distribution μ_3 with $\mu_2 \Rightarrow_c \mu_3$ and $\mu_3 \models \psi_2$ from which $s \models \langle\psi_1\langle\alpha\rangle\!\rangle\psi_2$ follows. $\square$

As a consequence of Lemma 20 extended in this manner, $\mathcal{L}^{state}_{\xi,\eta}{+}\epsilon$ is expressible in $\mathcal{L}^{state}_{\xi,b}{+}\epsilon$ and similarly $\mathcal{L}^{state}_{\xi,w}{+}\epsilon$ in $\mathcal{L}^{state}_{\xi,d}{+}\epsilon$. Theorem 2 would then refer to the following (partial) order

$$\begin{array}{lll} \nabla \preceq \Delta \preceq \lambda \preceq \varnothing & \nabla_{>0} \preceq \Delta_{>0} \preceq \mathsf{s} \preceq 0 \preceq \varnothing & \Delta \preceq \Delta_{>0} \\ b \preceq \eta \preceq w \quad b \preceq d \preceq w & & \end{array}$$

where, notably, $\nabla \not\preceq \nabla_{>0}$, and it would assert the following to hold:

- Any formula in $\mathcal{L}^{state}_{x_2,\xi_2}$ can be expressed in $\mathcal{L}^{state}_{x_1,\xi_1}$ whenever $\xi_1 \preceq \xi_2$ and $x_1 \preceq x_2$, except if $x_1 \in \{b, d\}$ and $x_2 \in \{\eta, w\}$.
- Any formula in $\mathcal{L}^{state}_{x_2,\xi_2}$ and $\mathcal{L}^{state}_{x_2,\xi_2}{+}\epsilon$ can be expressed in $\mathcal{L}^{state}_{x_1,\xi_1}{+}\epsilon$ whenever $\xi_1 \preceq \xi_2$ and $x_1 \preceq x_2$.

Recall that the combinations of $\nabla_{>0}$ with d or w are not considered, since no extended logic is defined.

9 Conclusion

This paper has explored the probabilistic bisimulation spectrum with silent moves both from an operational and from a logical perspective. In doing so, we have extended the arguably most popular fragment of van Glabbeek's linear-time branching-time spectrum with silent moves to the probabilistic setting. Since the extension is conservative, many of the results resemble those of van Glabbeek naturally, despite the fact that the probabilistic setting asks for some virtuosity in getting the proofs in place. Furthermore, the divergence dimension of the spectrum is more refined owed to the remarkable difference beween almost sure divergence and divergence with positive probability.

We have restricted ourselves to finite-state MDPs in the setup, but have been careful in flagging all results where the presented proofs carry over to countable MDPs. In effect, all of the spectrum can be rolled out for countable MDPs by (i) extending the notion of compound transitions to countable convex combinations (in countably branching τ-trees), mirrored on the logic side by (ii) moving from binary to countable conjunction (at the price of loosing denumerability). Only the results established for divergence probabilities and the logical characterisations for $\nabla_{>0}$-respecting branching and η-bisimilarity can not be extended in this manner, as we discussed. Indeed, the treatment of divergence probablities in $\Delta_{>0}$ and $\nabla_{>0}$ turned out rather challenging, not only because the reasons that led us to establish the needed results are very different for $\Delta_{>0}$ and $\nabla_{>0}$. Furthermore, we left the existence of a modal characterisation open for $\nabla_{>0}$ if combined with weak or delay bisimilarity.

References

1. Baier, C., Katoen, J.P.: Principles of Model Checking. MIT Press, Cambridge (2008)
2. Bergstra, J.A., Klop, J.W., Olderog, E.: Failures without chaos: a new process semantics for fair abstraction. In: Wirsing, M. (ed.) Formal Description of Programming Concepts-III: Proceedings of the IFIP TC 2/WG 2.2 Working Conference on Formal Description of Programming Concepts-III, Ebberup, Denmark, 25–28 Aug 1986, pp. 77–104. North-Holland (1987)
3. Blute, R., Desharnais, J., Edalat, A., Panangaden, P.: Bisimulation for labelled Markov processes. In: 12th Annual IEEE Symposium on Logic in Computer Science (LICS), pp. 149–158. IEEE Computer Society (1997)
4. Brengel, M.: Probabilistic weak transitions. Bachelor's thesis, Universität des Saarlandes, Fachrichtung Informatik (2013)
5. D'Argenio, P.R., Sánchez Terraf, P., Wolovick, N.: Bisimulations for non-deterministic labelled Markov processes. Math. Struct. Comput. Sci. **22**(1), 43–68 (2012). https://doi.org/10.1017/S0960129511000454
6. de Alfaro, L.: Formal verification of probabilistic systems. Ph.D. thesis, Stanford University, Department of Computer Science (1997)
7. Deng, Y., van Glabbeek, R.J., Hennessy, M., Morgan, C.: Testing finitary probabilistic processes. In: 20th International Conference on Concurrency Theory (CONCUR), Lecture Notes in Computer Science, vol. 5710, pp. 274–288. Springer, Berlin (2009)
8. Deng, Y., Hennessy, M.: On the semantics of Markov automata. Inf. Comput. **222**, 139–168 (2013). https://doi.org/10.1016/j.ic.2012.10.010
9. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Weak bisimulation is sound and complete for pCTL*. Inf. Comput. **208**(2), 203–219 (2010). https://doi.org/10.1016/j.ic.2009.11.002
10. Eisentraut, C., Hermanns, H., Zhang, L.: On probabilistic automata in continuous time. In: Proceedings of the 25th Annual IEEE Symposium on Logic in Computer Science, LICS 2010, 11–14 July 2010, Edinburgh, United Kingdom, pp. 342–351. IEEE Computer Society (2010). https://doi.org/10.1109/LICS.2010.41
11. Fischer, N., van Glabbeek, R.: Axiomatising infinitary probabilistic weak bisimilarity of finite-state behaviours. J. Log. Algebra Methods Program. **102**, 64–102 (2019). https://doi.org/10.1016/j.jlamp.2018.09.006
12. Groote, J.F., Vaandrager, F.W.: An efficient algorithm for branching bisimulation and stuttering equivalence. In: 17th International Colloquium on Automata, Languages and Programming (ICALP), Lecture Notes in Computer Science, vol. 443, pp. 626–638. Springer, Berlin (1990)
13. Hennessy, M.: Exploring probabilistic bisimulations, part I. Formal Asp. Comput. **24**(4–6), 749–768 (2012). https://doi.org/10.1007/s00165-012-0242-7
14. Hermanns, H., Parma, A., Segala, R., Wachter, B., Zhang, L.: Probabilistic logical characterization. Inf. Comput. **209**(2), 154–172 (2011). https://doi.org/10.1016/j.ic.2010.11.024
15. Lee, M.D., de Vink, E.P.: Logical characterization of bisimulation for transition relations over probability distributions with internal actions. In: Faliszewski, P., Muscholl, A., Niedermeier, R. (eds.) 41st International Symposium on Mathematical Foundations of Computer Science, MFCS 2016, 22–26 Aug 2016-Kraków, Poland, LIPIcs, vol. 58, pp. 29:1–29:14. Schloss Dagstuhl—Leibniz-Zentrum fuer Informatik (2016). https://doi.org/10.4230/LIPIcs.MFCS.2016.29
16. Liu, X., Yu, T., Zhang, W.: Logics for bisimulation and divergence. In: Baier, C., Lago, U.D. (eds.) Foundations of Software Science and Computation Structures—21st International Conference, FOSSACS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, 14–20 April 2018, Proceedings, Lecture Notes in Computer Science, vol. 10803, pp. 221–237. Springer, Berlin (2018). https://doi.org/10.1007/978-3-319-89366-2_12

17. Milner, R.: Communication and Concurrency. PHI Series in Computer Science. Prentice Hall, Upper Saddle River (1989)
18. Park, D.M.R.: Concurrency and automata on infinite sequences. In: Deussen, P. (ed.) Theoretical Computer Science, 5th GI-Conference, Karlsruhe, Germany, 23–25 March 1981, Proceedings, Lecture Notes in Computer Science, vol. 104, pp. 167–183. Springer, Berlin (1981). https://doi.org/10.1007/BFb0017309
19. Puterman, M.L.: Markov Decision Processes: Discrete Stochastic Dynamic Programming. Wiley, New York (1994)
20. Segala, R.: Modeling and verification of randomized distributed real-time systems. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, USA (1995). http://hdl.handle.net/1721.1/36560
21. Song, L., Zhang, L., Godskesen, J.C.: Bisimulations meet PCTL equivalences for probabilistic automata. Log. Methods Comput. Sci. (2013). https://doi.org/10.2168/LMCS-9(2:7)2013
22. Stoelinga, M.: Alea jacta est: verification of probabilistic, real-time and parametric systems. Ph.D. thesis, University of Nijmegen, the Netherlands (2002)
23. van Glabbeek, R.J.: The linear time-branching time spectrum II. In: Best, E. (ed.) CONCUR '93, 4th International Conference on Concurrency Theory, Hildesheim, Germany, 23–26 Aug 1993, Proceedings, Lecture Notes in Computer Science, vol. 715, pp. 66–81. Springer, Berlin (1993). https://doi.org/10.1007/3-540-57208-2_6. An extended version is available at http://boole.stanford.edu/pub/spectrum.pdf.gz
24. van Glabbeek, R.J.: The linear time-branching time spectrum I. In: Bergstra, J.A., Ponse, A., Smolka, S.A. (eds.) Handbook of Process Algebra, pp. 3–99. North-Holland, Elsevier, Amsterdam (2001). https://doi.org/10.1016/b978-044482830-9/50019-9
25. van Glabbeek, R.J., Luttik, B., Trcka, N.: Branching bisimilarity with explicit divergence. Fundam. Inform. **93**(4), 371–392 (2009). https://doi.org/10.3233/FI-2009-109
26. van Glabbeek, R.J., Weijland, W.P.: Branching time and abstraction in bisimulation semantics. J. ACM **43**(3), 555–600 (1996). https://doi.org/10.1145/233551.233556