Article 25fa pilot End User Agreement

This publication is distributed under the terms of Article 25fa of the Dutch Copyright Act (Auteurswet) with explicit consent by the author. Dutch law entitles the maker of a short scientific work funded either wholly or partially by Dutch public funds to make that work publicly available for no consideration following a reasonable period of time after the work was first published, provided that clear reference is made to the source of the first publication of the work.

This publication is distributed under The Association of Universities in the Netherlands (VSNU)'Article 25fa implementation' pilot project. In this pilot research outputs of researchers employed by Dutch Universities that comply with the legal requirements of Article 25fa of the Dutch Copyright Act are distributed online and free of cost or other barriers in institutional repositories. Research outputs are distributed six months after their first online publication in the original published version and with proper attribution to the source of the original publication.

You are permitted to download and use the publication for personal purposes. Please note that you are not allowed to share this article on other platforms, but can link to it. All rights remain with the author(s) and/or copyrights owner(s) of this work. Any use of the publication or parts of it other than authorised under this licence or copyright law is prohibited. Neither Radboud University nor the authors of this publication are liable for any damage resulting from your (re)use of this publication.

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please contact the Library through email: copyright@ubn.ru.nl, or send a letter to:

University Library
Radboud University
Copyright Information Point
PO Box 9100
6500 HA Nijmegen

You will be contacted as soon as possible.

(In)finite Trace Equivalence of Probabilistic Transition Systems

Alexandre Goy^{1,2} and Jurriaan Rot^{2(✉)}

¹ MICS - CentraleSupélec, Gif-sur-Yvette, France
alexandre.goy@student.ecp.fr

² Radboud University, Nijmegen, The Netherlands
j.rot@cs.ru.nl

Abstract. We show how finite and infinite trace semantics of generative probabilistic transition systems arises through a determinisation construction. This enables the use of bisimulations (up-to) to prove equivalence. In particular, it follows that trace equivalence for finite probabilistic transition systems is decidable. Further, the determinisation construction applies to both discrete and continuous probabilistic systems.

1 Introduction

The theory of coalgebras encompasses a wide variety of probabilistic systems, and according notions of bisimulation and behavioural equivalence [18]. We focus on one of the most basic instances: (generative) probabilistic transition systems (PTS), consisting of a set of states X and for every state a probability distribution over next states and (explicit) termination. Formally, they are coalgebras of the form $\alpha: X \rightarrow \mathcal{D}(A \times X + 1)$, where A is a fixed set of transition labels, $\mathcal{D}$ the probability distribution functor and $1 = \{*\}$ a singleton, whose element we interpret as an extra ‘accepting/termination’ state (Sect. 3 for details).

There is a natural notion of finite trace semantics for such PTSs, assigning to every state a sub-probability distribution of words, as a quantitative analogue of acceptance of words in non-deterministic automata. The definition of *infinite* traces is more subtle: it requires assigning probability to *sets* of traces rather than individual traces (infinite traces often have probability zero), and to move to probability *measures*. It is shown in [13] how finite and infinite trace semantics arises by modelling PTSs as coalgebras in the Kleisli category of the Giry monad.

As such, the (in)finite traces semantics of PTSs is an instance of the general theory of trace semantics through Kleisli categories, as proposed in [10]. A fundamentally different way of obtaining trace semantics of coalgebras is through

The research leading to these results has received funding from the European Research Council (FP7/2007–2013, grant agreement nr. 320571), and the Erasmus+ program. It was carried out during the first author’s master internship at Radboud University, reported in [9].

determinisation constructions, generalising the classical powerset construction of non-deterministic automata [12, 16] but also encompassing many other examples. In particular, in [12, 17] it is described how the *finite* traces of probabilistic transition systems arise through a certain determinisation construction, turning a PTS into a Moore automaton with sub-probability distributions as states. One of the advantages of determinisation is that it allows to use bisimulations (up-to) to prove trace equivalence. In particular, bisimulations up to congruence were used in Bonchi and Pous' HKC algorithm for non-deterministic automata [4] and in its extension to weighted automata [2].

In this paper, we show that the *finite and infinite* trace semantics of PTS, as in [13], arises through a determinisation construction (Sect. 4). The essential underlying idea that enables this approach, is that the (in)finite traces semantics in [13] is generated basically from two kinds of finite trace semantics: those that take into account termination/acceptance (as mentioned above), and those that do not (simply the probability of exhibiting a path in the PTS). In particular, for finite PTS, our determinisation construction yields an effective procedure for proving (in)finite trace equivalence using bisimulation up to congruence, using a variation of the HKC algorithm (Sect. 5). We finally show that the determinisation construction generalises to the setting of continuous PTS, working in the category of measurable spaces and with $\mathcal{D}$ replaced by the Giry monad (Sect. 6). While this generalises the discrete case, it is presented separately to make the discrete case accessible to a wider audience: the latter requires very little measure theory. We conclude with a discussion of related work (Sect. 7).

2 Preliminaries

Any finite set A can be called an *alphabet* and its elements *letters*. The set of words of length n with letters in A is denoted by A^n . By convention $A^0 = \{\varepsilon\}$ where ε is the empty word. The set of finite words over A is denoted by $A^* = \bigcup_{n \in \mathbb{N}} A^n$, the set of infinite words by $A^\omega = A^\mathbb{N}$ and the set of all (finite and infinite) words by $A^\infty = A^* \cup A^\omega$. A language L is a subset of $\mathcal{P}(A^*)$. It can be seen as a function $L: A^* \rightarrow \{0, 1\}$, by setting $L(w) = 1$ iff $w \in L$. The *language derivative* of L with respect to a letter a is defined by $L_a(w) = L(aw)$. The length of $w \in A^\infty$ is denoted by $|w| \in \mathbb{N} \cup \{\infty\}$. The concatenation function $c: A^* \times A^\infty \rightarrow A^\infty$ is denoted by juxtaposition ($c(u, v) = uv$) and defined by $uv(n) = u(n)$ if $n < |u|$ and $uv(n) = v(n - |u|)$ if $|u| \leq n < |u| + |v|$. It can be extended to languages $\mathcal{P}(A^*) \times \mathcal{P}(A^\infty) \rightarrow \mathcal{P}(A^\infty)$ by setting $LM = \{uv \mid u \in L, v \in M\}$. We sometimes abbreviate $\{w\}M$ by wM .

Coalgebras and Moore Automata. We recall the basic definition of coalgebras, see, e.g., [11, 15] for details and examples. The only instances that we use in this paper are Moore automata (recalled below), probabilistic transition systems (Sect. 3) and measure-theoretic generalisations of both (Sect. 6). Let $\mathcal{C}$ be a category, and $F: \mathcal{C} \rightarrow \mathcal{C}$ a functor. An F -coalgebra consists of an object X and an arrow $\alpha: X \rightarrow FX$. Given coalgebras (X, α) and (Y, β) , a coalgebra

homomorphism is an arrow $f: X \rightarrow Y$ such that $\beta \circ f = Ff \circ \alpha$. Coalgebras and homomorphisms form a category $\text{CoAlg}(F)$. A final object in $\text{CoAlg}(F)$ is called *final coalgebra*; explicitly, a coalgebra (Ω, ω) is final if for every F -coalgebra (X, α) there is unique coalgebra homomorphism $\varphi: X \rightarrow \Omega$. We recall the notion of *bisimulation* only for Moore automata, below.

Let B be a set. Define the *machine functor* $F_B: \text{Set} \rightarrow \text{Set}$ by $F_B X = B \times X^A$ and $Ff = \text{id}_B \times f^A$. An F_B -coalgebra $\langle o, t \rangle: X \rightarrow B \times X^A$ is called a *Moore automaton* (with output in B). A relation $R \subseteq X \times X$ on the states of a Moore automaton $\langle o, t \rangle: X \rightarrow B \times X^A$ is a *bisimulation* if for all $(x, y) \in R$: $o(x) = o(y)$ and for all $a \in A$, $(t_a(x), t_a(y)) \in R$ (here, we used the classical notation $t_a(x)$ instead of writing $t(x)(a)$). We write $x \sim y$ if there exists a bisimulation R such that xRy , and in this case say that x and y are *bisimilar*. For every B , there exists a final F_B -coalgebra (Ω, ω) where $\Omega = B^{A^*}$. For an F_B -coalgebra (X, α) , we write $\varphi_\alpha: X \rightarrow B^{A^*}$ or simply φ for the unique coalgebra morphism. We think of the elements of B^{A^*} as (weighted) languages, and of $\varphi(x)$ as the language of a state x . In particular, for $B = 2$, Moore automata are classical deterministic automata, and φ gives the usual language semantics. We have $\varphi(x) = \varphi(y)$ iff $x \sim y$, i.e., language equivalence coincides with bisimilarity.

Measure Theory. Let X be a set. A σ -algebra on X is a subset $\Sigma_X \subseteq \mathcal{P}(X)$ such that $\emptyset \in \Sigma_X$ and Σ_X is closed under complementation and countable union. Note that this implies that $X \in \Sigma_X$ and that Σ_X is closed under countable intersection and set difference. Given any subset $G \subseteq \mathcal{P}(X)$, there always exists a smallest σ -algebra containing G . Indeed, $\mathcal{P}(X)$ is a σ -algebra containing G , and the intersection of an arbitrary non-empty set of σ -algebras is itself a σ -algebra: just take the intersection of all σ -algebras containing G . We call it the σ -algebra *generated* by G and denote it by $\sigma_X(G)$. For example, $\mathcal{P}(X)$ is a σ -algebra on X . When working with real numbers $\mathbb{R}$, we will use the Borel σ -algebra $\mathcal{B}(\mathbb{R}) = \sigma_{\mathbb{R}}(\{(-\infty, x] \mid x \in \mathbb{R}\})$. We use $\mathcal{B}([0, 1]) = \{B \cap [0, 1] \mid B \in \mathcal{B}(\mathbb{R})\}$ as the canonical σ -algebra on $[0, 1]$. If X is a set and Σ_X is a σ -algebra on X , the pair (X, Σ_X) is called a measurable space. We write X for (X, Σ_X) when the σ -algebra used is clear. A function $f: (X, \Sigma_X) \rightarrow (Y, \Sigma_Y)$ is *measurable* if for all $S_Y \in \Sigma_Y$, $f^{-1}(S_Y) \in \Sigma_X$. The composition of measurable functions is measurable. An (implicitly finite) *measure* is a map $m: \Sigma_X \rightarrow \mathbb{R}_+$ such that $m(\emptyset) = 0$ and $m(\bigcup_{n \in \mathbb{N}} A_n) = \sum_{n \in \mathbb{N}} m(A_n)$ if the union is disjoint (σ -additivity property). We write $\mathcal{M}(X)$ for the set of measures on a measurable space (X, Σ_X) .

(Sub)distribution. The *distribution* functor $\mathcal{D}: \text{Set} \rightarrow \text{Set}$ is defined by $\mathcal{D}(X) = \{p: X \rightarrow [0, 1] \mid \sum_{x \in X} p(x) = 1\}$ and, given a function $f: X \rightarrow Y$, $\mathcal{D}f(u)(y) = \sum_{x \in f^{-1}(\{y\})} u(x)$. The functor $\mathcal{D}$ extends to a monad, with the unit η given by the Kronecker delta $\eta_X(x) = \delta_x$ (i.e., $\eta(x)(y) = 1$ if $x = y$ and $\eta(x)(y) = 0$ otherwise), and the multiplication by $\mu_X(U)(y) = \sum_{u \in \mathcal{D}X} U(u) \cdot u(y)$. The *sub-distribution* functor $\mathcal{S}: \text{Set} \rightarrow \text{Set}$ is defined by $\mathcal{S}(X) = \{p: X \rightarrow [0, 1] \mid \sum_{x \in X} p(x) \leq 1\}$. It extends to a monad in a similar way. There is a natural embedding of $\mathcal{D}$ in $\mathcal{S}$, which we denote by $\iota: \mathcal{D} \Rightarrow \mathcal{S}$.

3 Trace Semantics of Probabilistic Transition Systems

In this section, we recall PTSs and their (finite and infinite) trace semantics, following [13]. We start with the finite trace semantics.

Definition 3.1. A probabilistic transition system (PTS) is a coalgebra for the functor $\mathcal{D}(A \times \text{Id} + 1)$, i.e., a set X together with a map $\alpha: X \rightarrow \mathcal{D}(A \times X + 1)$.

Definition 3.2. Let $\alpha: X \rightarrow \mathcal{D}(A \times X + 1)$ be a PTS. The finite trace semantics $\llbracket - \rrbracket_f: X \rightarrow \mathcal{S}(A^*)$ is defined by the following equations.

$$\llbracket x \rrbracket_f(\varepsilon) = \alpha(x)(*) \quad \llbracket x \rrbracket_f(aw) = \sum_{y \in X} \alpha(x)(a, y) \cdot \llbracket y \rrbracket_f(w)$$

for all $x \in X$, $a \in A$, $w \in A^*$.

Consider as a first example the simple PTS below, where the element $*$ is represented as a distinguished double-circled state, and a transition is represented by an arrow labeled with its probability.

We have $\llbracket x \rrbracket_f(a^n) = \frac{1}{2^{n+1}}$ for all n . The trace semantics becomes more subtle if infinite words are also taken into account. Consider, for instance, the following PTS.

Intuitively both states accept any finite or infinite word w over $\{a, b\}$ with probability 0. However, the probability of ‘starting with an a ’ in y or z is clearly different. This becomes apparent when we move to assigning probability to sets of traces, which is where we need a bit of measure theory. We therefore first define a suitable σ -algebra on the set A^∞ of finite and infinite words.

Definition 3.3. Let $S_\infty = \{\emptyset\} \cup \{\{w\} \mid w \in A^*\} \cup \{wA^\infty \mid w \in A^*\}$. The σ -algebra of measurable sets of words is defined to be $\Sigma_{A^\infty} = \sigma_{A^\infty}(S_\infty)$.

This σ -algebra is generated by a countable family of generators: the empty set, the singletons of finite words, and the cones, i.e., sets wA^∞ of words that have the finite word w as a prefix. This σ -algebra is very natural. Indeed, the usual measure-theoretical σ -algebra on $A^* \cup A^\omega$ would be the combination of the discrete σ -algebra $\mathcal{P}(A^*)$ and the product σ -algebra (see, e.g., [1], Definition 4.42) of all $\mathcal{P}(A)$ on A^ω . One can easily prove that this construction yields our Σ_{A^∞} too. In the sequel, this is the σ -algebra on A^∞ implicitly used. The following proposition establishes measurability for some useful sets.

Proposition 3.4. The following sets of words are measurable:

- (i) The singleton $\{w\}$ for any $w \in A^\infty$;
- (ii) any countable language;
- (iii) any language of finite words;
- (iv) $\emptyset, A^*, A^\omega, A^\infty$;
- (v) the concatenation LS where $L \subseteq A^*$ and $S \in \Sigma_{A^\infty}$.

In the following, if m is a measure over A^∞ and $w \in A^\infty$, we will write $m(w)$ instead of $m(\{w\})$. We have the following key theorem, which follows easily from results in [13]:

Theorem 3.5. *Let $m: S_\infty \rightarrow \mathbb{R}_+$ be a map satisfying $m(\emptyset) = 0$. The two following conditions are equivalent.*

- (i) *There exists a unique measure $\tilde{m}: \Sigma_{A^\infty} \rightarrow \mathbb{R}_+$ such that $\tilde{m}|_{S_\infty} = m$.*
- (ii) *For all $w \in A^*$, $m(wA^\infty) = m(w) + \sum_{a \in A} m(waA^\infty)$.*

Proof. (i) $\Rightarrow$ (ii) The equation comes directly from the σ -additivity of $\tilde{m}$. (ii) $\Rightarrow$ (i) According to Lemma 3.18 of [13], (ii) is equivalent to the fact that m is a pre-measure. Using Caratheodory's extension theorem (e.g., [14]), this pre-measure can be uniquely extended to a measure as in (i). $\square$

Recall that $\mathcal{M}(A^\infty)$ denotes the set of measures m on A^∞ .

Definition 3.6. *Let $\alpha: X \rightarrow \mathcal{D}(A \times X + 1)$ be a PTS. The (finite and infinite) trace semantics $\llbracket - \rrbracket: X \rightarrow \mathcal{M}(A^\infty)$ is defined by the following equations.*

$$\begin{aligned} \llbracket x \rrbracket(\varepsilon A^\infty) &= 1 & \llbracket x \rrbracket(\varepsilon) &= \alpha(x)(*) \\ \llbracket x \rrbracket(awA^\infty) &= \sum_{y \in X} \alpha(x)(a, y) \cdot \llbracket y \rrbracket(wA^\infty) & \llbracket x \rrbracket(aw) &= \sum_{y \in X} \alpha(x)(a, y) \cdot \llbracket y \rrbracket(w) \end{aligned}$$

(These equations uniquely determine a measure by Theorem 3.5.)

Example 3.7. Consider the following PTS over the alphabet $A = \{a\}$.

The semantics $\llbracket x \rrbracket$ is easy to compute for sets of words in S_∞ by induction: for every finite word w , $\llbracket x \rrbracket(w) = 0$ and $\llbracket x \rrbracket(wA^\infty) = 1$. Because $a^n A^\infty$ is a non-increasing sequence of measurable sets converging to $\{a^\omega\}$, properties of measures yield $\llbracket x \rrbracket(a^\omega) = \lim_{n \rightarrow +\infty} \llbracket x \rrbracket(a^n A^\infty) = 1$. Let us look at $\llbracket y \rrbracket$. Intuitively, the probability of performing n loops in state y and then moving to (and staying in) state x is $1/3^{n+1}$. Summing them for $n \in \mathbb{N} \cup \{0\}$ gives $1/2$, the

probability of moving eventually to state x . Indeed, first observe that $\llbracket y \rrbracket(\varepsilon) = 1/3$ and $\llbracket y \rrbracket(\varepsilon A^\infty) = 1$. Let $n \in \mathbb{N} \cup \{0\}$, then:

$$\begin{aligned}\llbracket y \rrbracket(a^{n+1}) &= \frac{1}{3}\llbracket y \rrbracket(a^n) + \frac{1}{3}\llbracket x \rrbracket(a^n) = \frac{1}{3}\llbracket y \rrbracket(a^n) \\ \llbracket y \rrbracket(a^{n+1}A^\infty) &= \frac{1}{3}\llbracket y \rrbracket(a^n A^\infty) + \frac{1}{3}\llbracket x \rrbracket(a^n A^\infty) = \frac{1}{3}\llbracket y \rrbracket(a^n A^\infty) + \frac{1}{3}\end{aligned}$$

One can then prove that $\llbracket y \rrbracket(a^n) = 1/3^{n+1}$ and $\llbracket y \rrbracket(a^\omega) = \lim_{n \rightarrow +\infty} \llbracket y \rrbracket(a^n A^\infty) = \lim_{n \rightarrow +\infty} (1 + 3^{-n})/2 = 1/2$.

Example 3.8. Consider again the PTS in (2). We have $\llbracket y \rrbracket(w) = \llbracket z \rrbracket(w) = 0$ for all $w \in A^*$. However, $\llbracket y \rrbracket(aA^\infty) = \llbracket y \rrbracket(bA^\infty) = 1/2$ whereas $\llbracket z \rrbracket(aA^\infty) = 3/4$ and $\llbracket z \rrbracket(bA^\infty) = 1/4$. Hence $\llbracket y \rrbracket \neq \llbracket z \rrbracket$, as expected.

The above (in)finite trace semantics is essentially generated from two kinds of finite trace semantics: one for finite words w , and one for cones wA^∞ , where $w \in A^*$. The probability of the latter is simply the probability the finite trace w without considering acceptance/termination, i.e., the probability of exhibiting the path w . This finite presentation is exploited in the determinisation construction in the next section, which essentially encodes both kinds of finite trace semantics simultaneously.

4 Determinisation

In this section, we show how the finite and infinite trace semantics of PTS (Definition 3.6) arises through a determinisation construction. This construction transforms any PTS into a certain kind of Moore machine with sub-probability distributions as states. The final coalgebra semantics of this Moore machine represents the trace semantics $\llbracket - \rrbracket: X \rightarrow \mathcal{M}(A^\infty)$ of the original PTS. The determinisation procedure is exploited in the next section to give an algorithm for computing (in)finite trace equivalence, based on bisimulations.

In Sect. 6, we consider a more general kind of PTS, with measurable sets as state spaces, which fully generalises the results and constructions of the current section. Most proofs in the current section are hence omitted. Moreover, it is explained in Sect. 6 that our approach is an instance of the abstract framework of coalgebraic determinisation based on distributive laws [12, 16]. In the current section we mostly neglect this and present the concrete constructions.

Throughout this section, let $\alpha: X \rightarrow \mathcal{D}(A \times X + 1)$ be a PTS. Our approach to (in)finite traces resembles the determinisation construction of [12, 17] for finite traces of PTSs. As explained below, there is one crucial addition for (in)finite traces: we make the total weight of sub-distributions in the determinised coalgebra observable, essentially to capture the probability of the ‘cones’ wA^∞ . We will show that the resulting final coalgebra semantics factorises through the set $\mathcal{M}(A^\infty)$ of measures on words, recovering the trace semantics of Definition 3.6. The overall construction is as follows.

- (i) Translate α into a coalgebra $\tilde{\alpha}: X \rightarrow [0, 1] \times [0, 1] \times (\mathcal{S}X)^A$.
- (ii) Determinise it: define an $\alpha^\sharp: \mathcal{S}X \rightarrow [0, 1] \times [0, 1] \times (\mathcal{S}X)^A$ such that $\alpha^\sharp \circ \eta_X = \tilde{\alpha}$. Let $\varphi: \mathcal{S}X \rightarrow ([0, 1] \times [0, 1])^{A^*}$ be the unique map to the final coalgebra.
- (iii) Factorise φ to get a coalgebra morphism $\mathcal{S}X \rightarrow \mathcal{M}(A^\infty)$, then precompose with η_X to get the desired trace semantics $X \rightarrow \mathcal{M}(A^\infty)$.

The construction is summed up in the following diagram. Below, we explain each of the steps in detail.

$$\begin{array}{ccccccc}
 X & \xrightarrow{\eta_X} & \mathcal{S}X & \xrightarrow{\varphi} & \mathcal{M}(A^\infty) & \xrightarrow{\quad} & ([0, 1] \times [0, 1])^{A^*} \\
 \alpha \downarrow & \searrow \tilde{\alpha} & \downarrow \alpha^\sharp & \nearrow [-] & \downarrow \Pi & & \downarrow \omega \\
 \mathcal{D}(A \times X + 1) & & [0, 1] \times [0, 1] \times (\mathcal{S}X)^A & \xrightarrow{\quad} & [0, 1] \times [0, 1] \times (\mathcal{M}(A^\infty))^A & \xrightarrow{\quad} & [0, 1] \times [0, 1] \times (([0, 1] \times [0, 1])^{A^*})^A \\
 \downarrow \iota_{A \times X + 1} & \searrow \epsilon_X & \downarrow & \nearrow \text{id} \times \varphi^A & \downarrow & & \downarrow \\
 \mathcal{S}(A \times X + 1) & \xrightarrow{\quad} & [0, 1] \times [0, 1] \times (\mathcal{S}X)^A & \xrightarrow{\quad} & [0, 1] \times [0, 1] \times (\mathcal{M}(A^\infty))^A & \xrightarrow{\quad} & [0, 1] \times [0, 1] \times (([0, 1] \times [0, 1])^{A^*})^A
 \end{array}$$

Remark 4.1. As mentioned above, the construction is quite close to the determinisation of finite traces [12, 17]. There are two main differences: first, the latter determinises to a Moore automaton of the type $\mathcal{S}X \rightarrow [0, 1] \times (\mathcal{S}X)^A$ (so with $[0, 1]$ rather than $[0, 1] \times [0, 1]$). Second, the decomposition of φ here yields measures (to represent (in)finite trace semantics) rather than sub-probability distributions (to represent finite trace semantics).

(i) Translation: from α to $\tilde{\alpha}$. The first step, the definition of $\tilde{\alpha}$ from α , basically forgets certain information about probability distributions. The natural transformation ϵ is given on a component X by $\epsilon_X: \mathcal{S}(A \times X + 1) \rightarrow [0, 1] \times [0, 1] \times (\mathcal{S}X)^A$,

$$\epsilon_X(u) = \left\langle \sum_{z \in A \times X + 1} u(z), u(*), a \mapsto [y \mapsto u(a, y)] \right\rangle.$$

We have $\tilde{\alpha}(x) = \langle 1, \alpha(x)(*), a \mapsto [y \mapsto \alpha(x)(a, y)] \rangle$.

(ii) Determinisation. In the second step, we turn $\tilde{\alpha}$ into a Moore automaton over sub-distributions. Formally, the latter will be a coalgebra for the functor $F_{[0,1] \times [0,1]}: \mathbf{Set} \rightarrow \mathbf{Set}$; recall from Sect. 2 that this is defined by $F_{[0,1] \times [0,1]}X = [0, 1] \times [0, 1] \times X^A$. In the remainder of this section we abbreviate $F_{[0,1] \times [0,1]}$ by F . Notice that $\tilde{\alpha}$ is an $F\mathcal{S}$ -coalgebra. Any $F\mathcal{S}$ -coalgebra determinises to an F -coalgebra, but we spell it out here only for the necessary instance $\tilde{\alpha}$. For a concrete example, see the first part of Example 5.9 in the next section.

Definition 4.2. *The determinisation of the PTS $\alpha: X \rightarrow \mathcal{D}(A \times X + 1)$ is the Moore machine $\alpha^\sharp: \mathcal{S}X \rightarrow F\mathcal{S}X = [0, 1] \times [0, 1] \times (\mathcal{S}X)^A$, defined by:*

$$\alpha^\sharp(u) = \left\langle \sum_{x \in X} u(x), \sum_{x \in X} u(x) \cdot \alpha(x)(*), a \mapsto [y \mapsto \sum_{x \in X} u(x) \cdot \alpha(x)(a, y)] \right\rangle.$$

(iii) Factorisation of final coalgebra semantics. Since $\alpha^\sharp$ is an F -coalgebra, there exists a unique coalgebra morphism φ from $(SX, \alpha^\sharp)$ to the final coalgebra $(([0, 1] \times [0, 1])^{A^*}, \omega)$. This is not quite the right type: the (in)finite trace semantics $\llbracket - \rrbracket$ is a (probability) measure over words, hence, for each x , $\llbracket x \rrbracket$ should be an element of the set $\mathcal{M}(A^\infty)$ of measures over A^∞ (Sect. 2). In the last step (iii), we equip $\mathcal{M}(A^\infty)$ with an F -coalgebra structure Π which is final among F -coalgebras satisfying a certain property, satisfied by our determinisation $\alpha^\sharp$. This allows us to factor φ through a coalgebra homomorphism $[-]: \mathcal{S}(X) \rightarrow \mathcal{M}(A^\infty)$. In the more general setting of Sect. 6 we show how the coalgebra structure on $\mathcal{M}(A^\infty)$ arises from the Giry monad and the final coalgebra of the **Set** endofunctor $X \mapsto A \times X + 1$. For now, we define it explicitly, which requires:

Definition 4.3 (Measure derivative). *Let m be a measure on A^∞ and $a \in A$. The map m_a defined by $m_a(S) = m(aS)$ for any $S \in \Sigma_{A^\infty}$ is a measure, called the measure derivative of m (with respect to a).*

It is easy to check that m_a as defined above is indeed a measure, so that the measure derivative is well-defined. Now, the coalgebra $\Pi: \mathcal{M}(A^\infty) \rightarrow [0, 1] \times [0, 1] \times (\mathcal{M}(A^\infty))^A$ is defined by $\Pi: m \mapsto \langle m(\varepsilon A^\infty), m(\varepsilon), a \mapsto m_a \rangle$. Since Π is an F -coalgebra, we obtain a coalgebra morphism to the final F -coalgebra ω .

Lemma 4.4. *The unique coalgebra morphism from Π to ω is injective.*

A proof is given in the more general setting of Lemma 6.8. The following crucial lemma states in which cases the factorisation is possible. It establishes the F -coalgebra Π as a final object in a certain subcategory of $\text{CoAlg}(F)$.

Proposition 4.5. *Let $\beta = \langle \beta_\oplus, \beta_*, a \mapsto \tau_a \rangle: Y \rightarrow FY$ be an F -coalgebra. The two following conditions are equivalent:*

- (i) *There exists an F -coalgebra morphism $[-]$ from β to Π .*
- (ii) *The equation $\beta_\oplus = \beta_* + \sum_{a \in A} \beta_\oplus \circ \tau_a$ holds.*

In this case, this morphism is unique.

See Theorem 6.9 for a proof in the (more general) continuous setting.

Lemma 4.6. *The coalgebra $\alpha^\sharp: SX \rightarrow [0, 1] \times [0, 1] \times (SX)^A$ satisfies (ii) in Proposition 4.5.*

It is important to note that condition (ii) does not hold in general if the whole construction starts from a coalgebra of the form $\alpha: X \rightarrow \mathcal{S}(A \times X + 1)$. The price to be paid for a PTS to be compatible enough to generate infinite trace semantics from the finite traces in a measure-theoretic way is to sum to 1, i.e., to use $\mathcal{D}$ and not $\mathcal{S}$.

The following result summarises the situation.

Corollary 4.7. *The morphism φ decomposes as a unique coalgebra morphism $[-]: SX \rightarrow \mathcal{M}(A^\infty)$ from $\alpha^\sharp$ to Π followed by an injective coalgebra morphism φ_Π from Π to ω , as shown in the following diagram.*

$$\begin{array}{ccccc}
 SX & \xrightarrow{\quad [-] \quad} & \mathcal{M}(A^\infty) & \xrightarrow{\quad \varphi_\Pi \quad} & ([0, 1] \times [0, 1])^{A^*} \\
 \alpha^\sharp \downarrow & & \downarrow \Pi & & \downarrow \omega \\
 [0, 1] \times [0, 1] \times (SX)^A & \longrightarrow & [0, 1] \times [0, 1] \times (\mathcal{M}(A^\infty))^A & \longrightarrow & [0, 1] \times [0, 1] \times ([0, 1] \times [0, 1])^{A^*}
 \end{array}$$

φ

We thus obtain the semantics $[-] \circ \eta_X: X \rightarrow \mathcal{M}(A^\infty)$ by precomposing with the unit of the monad $\mathcal{S}$. It coincides with the semantics $\llbracket - \rrbracket$ of Definition 3.6:

Theorem 4.8. *We have $\llbracket - \rrbracket = [-] \circ \eta_X$.*

Theorem 4.8 is the main result of this section, stating that the (in)finite trace semantics is recovered by finality through a determinisation construction. Together with Lemma 4.4, it yields equivalence between the first two points below.

Corollary 4.9. *For any $x, y \in X$, the following are equivalent:*

1. $\llbracket x \rrbracket = \llbracket y \rrbracket$,
2. $\varphi(\delta_x) = \varphi(\delta_y)$,
3. $\delta_x \sim \delta_y$,

where $\sim$ is bisimilarity on the Moore automaton $\alpha^\sharp$, the determinisation of α (Definition 4.2).

The equivalence between 2. and 3. is standard, and was mentioned in Sect. 2. By the equivalence between 1. and 3., we can prove (in)finite trace equivalence by computing bisimulations, which is used in the next section.

5 Computing Trace Equivalence

The aim of this section is to give an algorithm that takes states $x, y \in X$ of a PTS and tells whether x and y are (in)finite trace equivalent (i.e., $\llbracket x \rrbracket = \llbracket y \rrbracket$) or not, based on the determinisation construction described in Sect. 4. Our algorithm is a variant of HKC, an algorithm for language equivalence of non-deterministic automata based on determinisation and bisimulation (up-to) techniques [4]. More specifically, we will use its generalisation to weighted automata given in [2].

Let $\alpha: X \rightarrow \mathcal{D}(A \times X + 1)$ be a *finite-state* PTS and $\alpha^\sharp$ its determinisation (Definition 4.2). By Corollary 4.9, to prove $\llbracket x \rrbracket = \llbracket y \rrbracket$ it suffices to show $\delta_x \sim \delta_y$, i.e., that there is a bisimulation $R \subseteq SX \times SX$ on the determinised Moore automaton such that $(\delta_x, \delta_y) \in R$. However, this task can be simplified using *bisimulation up-to techniques*, as explained next. In order to use the techniques from [2], we first move from sub-probability distributions to vector spaces. To this

end, define $\mathbb{R}_\omega^X$ as the set of finitely supported functions $X \rightarrow \mathbb{R}$, i.e., $\mathbb{R}_\omega^X = \{u: X \rightarrow \mathbb{R} \mid u(x) \neq 0 \text{ for finitely many } x\}$. We define the Moore automaton $\bar{\alpha} = \langle \bar{\alpha}_\oplus, \bar{\alpha}_*, a \mapsto \bar{\alpha}_a \rangle: \mathbb{R}_\omega^X \rightarrow \mathbb{R} \times \mathbb{R} \times (\mathbb{R}_\omega^X)^A$ as follows on any $u \in \mathbb{R}_\omega^X$:

$$\bar{\alpha}_\oplus = \sum_{x \in X} u(x) \quad \bar{\alpha}_* = \sum_{x \in X} u(x) \cdot \alpha(x)(*) \quad \bar{\alpha}_a = \left[y \mapsto \sum_{x \in X} u(x) \cdot \alpha(x)(a, y) \right] \quad (3)$$

This is almost the same construction as in Definition 4.2, with sub-probability distributions replaced by vectors. (Note that this is well-defined since X is assumed to be finite; it would also suffice to assume that α is finitely branching.) It is easy to see that the embedding $i: \mathcal{S}X \rightarrow \mathbb{R}_\omega^X$ is an injective $F_{[0,1] \times [0,1]}$ -coalgebra morphism from $\alpha^\sharp$ to $\bar{\alpha}$. Together with Corollary 4.9, this yields:

Corollary 5.1. *For any $x, y \in X$: $\llbracket x \rrbracket = \llbracket y \rrbracket$ iff $\delta_x \sim \delta_y$, where $\sim$ is bisimilarity on the Moore automaton $\bar{\alpha}$.*

We now formulate bisimulation up to congruence, concretely for $\bar{\alpha}$.

Definition 5.2. *Let $R \subseteq \mathbb{R}_\omega^X \times \mathbb{R}_\omega^X$. Its congruence closure $c(R)$ is the least congruence that contains R , i.e., that satisfies*

$$\frac{(u, v) \in R}{(u, v) \in c(R)} \quad \frac{}{(u, u) \in c(R)} \quad \frac{(u, v) \in c(R)}{(v, u) \in c(R)} \quad \frac{(u, v) \in c(R) \quad (v, w) \in c(R)}{(u, w) \in c(R)}$$

$$\frac{(r \cdot u, r \cdot v) \in c(R)}{(u, v) \in c(R)} (r \in \mathbb{R}) \quad \frac{(u, u') \in c(R) \quad (v, v') \in c(R)}{(u + u', v + v') \in c(R)}$$

Definition 5.3. *Define $\bar{\alpha}: \mathbb{R}_\omega^X \rightarrow \mathbb{R} \times \mathbb{R} \times (\mathbb{R}_\omega^X)^A$ from a finite-state PTS α , as in Eq. (3). A relation $R \subseteq \mathbb{R}_\omega^X \times \mathbb{R}_\omega^X$ is a bisimulation up to congruence (on $\bar{\alpha}$) if for all $(u, v) \in R$:*

- $\bar{\alpha}_\oplus(u) = \bar{\alpha}_\oplus(v)$, $\bar{\alpha}_*(u) = \bar{\alpha}_*(v)$, and
- $\forall a \in A: (\bar{\alpha}_a(u), \bar{\alpha}_a(v)) \in c(R)$.

The following result states soundness of bisimulations up to congruence. This can either be proved from the abstract coalgebraic theory [3] or more directly using compatible functions, as in [2, 4].

Theorem 5.4. *For any $u, v \in \mathbb{R}_\omega^X$: $u \sim v$ iff there is a bisimulation up to congruence R (on $\bar{\alpha}$) such that $(u, v) \in R$.*

Combined with Corollary 5.1, this means that to prove that $\llbracket x \rrbracket = \llbracket y \rrbracket$ for states x, y of a PTS, it suffices to show that there is a bisimulation up to congruence relating δ_x and δ_y . The following algorithm attempts to compute one given x, y .

$$\text{HKC}^\infty(x, y)$$

```

(1)  $R := \emptyset$ ;  $todo := \emptyset$ 
(2) insert  $(\delta_x, \delta_y)$  into  $todo$ 
(3) while  $todo$  is not empty do
  (3.1) extract  $(u, v)$  from  $todo$ 
  (3.2) if  $(u, v) \in c(R)$  then continue
  (3.3) if  $\bar{\alpha}_\oplus(u) \neq \bar{\alpha}_\oplus(v)$  then return false
  (3.3') if  $\bar{\alpha}_*(u) \neq \bar{\alpha}_*(v)$  then return false
  (3.4) for all  $a \in A$ , insert  $(\bar{\alpha}_a(u), \bar{\alpha}_a(v))$  into  $todo$ 
  (3.5) insert  $(u, v)$  into  $R$ 
(4) return true

```

Theorem 5.5. *Whenever $\text{HKC}^\infty(x, y)$ terminates, it returns **true** iff $\llbracket x \rrbracket = \llbracket y \rrbracket$.*

Despite the fact that during the determinisation the state space always becomes infinite, the following results show that if the initial state space X is finite, then HKC^∞ does terminate.

Theorem 5.6 (see [6]). *Let $\mathcal{R}$ be a ring and X be a finite set. Let $R \subseteq \mathcal{R}^X \times \mathcal{R}^X$ be a relation and let $(v, v') \in \mathcal{R}^X \times \mathcal{R}^X$ be a pair of vectors. Let $U_R = \{u - u' \mid (u, u') \in R\}$. Then $(v, v') \in c(R)$ iff $v - v' \in [U_R]$, where $[U_R]$ is the submodule of $\mathcal{R}^X$ generated by U_R .*

Proposition 5.7. *If X is finite, $\text{HKC}^\infty(x, y)$ terminates for every $x, y \in X$.*

Example 5.8. To begin with, here is a very simple PTS which we use to demonstrate the need for bisimulation up to congruence over plain bisimulations.

A bisimulation on the determinised automaton containing (δ_x, δ_y) would require adding $(\delta_x/2^k, \delta_y/2^k)$ for all k to the relation. However, $\text{HKC}^\infty(x, y)$ (which computes a bisimulation up to congruence) stops after one step because it spots that $(\delta_x/2, \delta_y/2)$ is in the congruence closure of the relation $\{(\delta_x, \delta_y)\}$.

Example 5.9. Consider the PTS depicted on the left below. We will use HKC^∞ to check if the states x and z are (in)finite trace equivalent.

$$u = \begin{pmatrix} u_x \\ u_y \\ u_z \\ u_i \end{pmatrix} \quad M_a = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 1/6 & 1/3 & 0 & 0 \\ 0 & 0 & 1/3 & 0 \\ 1/2 & 0 & 1/3 & 1 \end{pmatrix}$$

$$L = \begin{pmatrix} L_\oplus \\ L_* \end{pmatrix}$$

First, we compute part of the determinised automaton. To this end, observe that because X is finite, $\mathbb{R}_\omega^X = \mathbb{R}^X$ has a basis (e_x, e_y, e_z, e_i) . An element $u \in \mathbb{R}_\omega^X$ is seen as a column vector $u_x e_x + u_y e_y + u_z e_z + u_i e_i$ in this basis. Moreover $\bar{\alpha}_\oplus$ and $\bar{\alpha}_*$ are linear forms that can be seen as the row vectors $L_\oplus = (1 \ 1 \ 1 \ 1)$ and $L_* = (1/3 \ 2/3 \ 1/3 \ 0)$, and $\bar{\alpha}_a$ is an endomorphism with a transition matrix M_a defined by $(M_a)_{j,k} = t_a(k)(j)$. This is depicted on the right above.

We represent here two parts of the determinised automaton. The first is the path beginning with the single state x ; the second is the path beginning with the single state z . Each state here has two real outputs, obtained by matrix multiplication with $L_\oplus$ and L_* .

$$\begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} \xrightarrow{a} \begin{pmatrix} 0 \\ 1/6 \\ 0 \\ 1/2 \end{pmatrix} \xrightarrow{a} \begin{pmatrix} 0 \\ 1/18 \\ 0 \\ 1/2 \end{pmatrix} \xrightarrow{a} \dots \xrightarrow{a} \begin{pmatrix} 0 \\ 1/(2 \times 3^n) \\ 0 \\ 1/2 \end{pmatrix} \xrightarrow{a} \dots$$

$$\vdots \quad \vdots \quad \vdots \quad \vdots$$

$$\begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} \xrightarrow{a} \begin{pmatrix} 0 \\ 0 \\ 1/3 \\ 1/3 \end{pmatrix} \xrightarrow{a} \begin{pmatrix} 0 \\ 0 \\ 1/9 \\ 4/9 \end{pmatrix} \xrightarrow{a} \dots \xrightarrow{a} \begin{pmatrix} 0 \\ 0 \\ 1/3^n \\ (1 - 3^{-n})/2 \end{pmatrix} \xrightarrow{a} \dots$$

Now, $\text{HKC}^\infty(x, z)$ begins with $\text{todo} = \{(\eta_X(x), \eta_X(z))\} = \{(e_x, e_z)\}$ and $R = \emptyset$. It checks that $Le_x = Le_z$, etc. as shown in the following table.

The check succeeds in loop 3 because $(u, v) \in c(R)$ according to Theorem 5.6:

$$\begin{pmatrix} 0 \\ 1/18 \\ 0 \\ 1/2 \end{pmatrix} - \begin{pmatrix} 0 \\ 0 \\ 1/9 \\ 4/9 \end{pmatrix} = \begin{pmatrix} 0 \\ 1/18 \\ -1/9 \\ 1/18 \end{pmatrix} = \frac{1}{3} \begin{pmatrix} 0 \\ 1/6 \\ -1/3 \\ 1/6 \end{pmatrix} = \frac{1}{3} \left(\begin{pmatrix} 0 \\ 1/6 \\ 0 \\ 1/2 \end{pmatrix} - \begin{pmatrix} 0 \\ 0 \\ 1/3 \\ 1/3 \end{pmatrix} \right)$$

Because **todo** is eventually empty, the algorithm returns **true**. Indeed, if we compute directly the measures $\llbracket x \rrbracket$ and $\llbracket z \rrbracket$, we can see that $\llbracket x \rrbracket(a^n) = 1/3^{n+1}$, $\llbracket x \rrbracket(a^\omega) = 1/2$ and similarly for $\llbracket z \rrbracket$. Here the bisimulation up to congruence check is necessary for termination. The construction of a bisimulation up to equivalence (dashed + dotted lines on the determinised automaton picture) would take an infinite number of steps. But the construction of the bisimulation up to congruence (dashed lines) takes only 2 steps.

Step	(3.1)	(3.2)	(3.3)	(3.4)	(3.5)
Loop counter	(u, v) extracted from todo	Check $(u, v) \in c(R)$	Check $Lu = Lv$	$(M_a u, M_a v)$ added to todo	Cardinality of R
1	$\left(\begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} \right)$	Fail	$\begin{pmatrix} 1 \\ 1/3 \end{pmatrix} = \begin{pmatrix} 1 \\ 1/3 \end{pmatrix}$	$\left(\begin{pmatrix} 0 \\ 1/6 \\ 0 \\ 1/2 \end{pmatrix}, \begin{pmatrix} 0 \\ 1/3 \\ 1/3 \end{pmatrix} \right)$	1
2	$\left(\begin{pmatrix} 0 \\ 1/6 \\ 0 \\ 1/2 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1/3 \\ 1/3 \end{pmatrix} \right)$	Fail	$\begin{pmatrix} 2/3 \\ 1/9 \end{pmatrix} = \begin{pmatrix} 2/3 \\ 1/9 \end{pmatrix}$	$\left(\begin{pmatrix} 0 \\ 1/18 \\ 0 \\ 1/2 \end{pmatrix}, \begin{pmatrix} 0 \\ 1/9 \\ 1/9 \\ 4/9 \end{pmatrix} \right)$	2
3	$\left(\begin{pmatrix} 0 \\ 1/18 \\ 0 \\ 1/2 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1/9 \\ 4/9 \end{pmatrix} \right)$	Success	/	/	2
4	Empty	/	/	/	/

6 Continuous Systems

In this section, we generalise the determinisation construction for (in)finite trace semantics previously defined to the case of continuous PTS, defined later as coalgebras for the analogue of functor $\mathcal{D}(A \times - + 1)$ in the category **Meas** (see [13] for examples of such PTSs). The underlying distributive law is brought to light, so that the origin of the determinisation process is better understood. The following table sums up the analogies and differences with the discrete case.

	Discrete case	General case
Category	Set	Meas
Usual operation	$\sum$	$\int$
Machine functor	$FX = [0, 1] \times [0, 1] \times X^A$	Measurable version of F
Probability monad	Distribution monad $\mathcal{D}$	Giry's monad $\mathbb{D}$
Determinisation monad	Sub-distribution monad $\mathcal{S}$	Sub-Giry's monad $\mathbb{S}$
PTS state space	Set X	Measurable space (X, Σ_X)
Determinised state	Finitely supported vector	Measure (≤ 1)
Transitions	Matrix $t_a: X \times X \rightarrow [0, 1]$	Kernel $t_a: X \times \Sigma_X \rightarrow [0, 1]$
Final F -coalgebra	ω	Measurable version of ω
Measure coalgebra	Π	Measurable version of Π
Pseudo-final morphism	$[-]: \mathcal{S}X \rightarrow \mathcal{M}(A^\infty)$	$[-]: \mathbb{S}X \rightarrow \mathbb{S}A^\infty$

In this section we work in the category **Meas** of measurable spaces and functions. It is easy to adapt F , but considering the monads we will need some additional measure-theoretic background.

Product. Given measurable spaces (X, Σ_X) and (Y, Σ_Y) , we define a product σ -algebra on $X \times Y$ by $\Sigma_X \otimes \Sigma_Y = \sigma_{X \times Y}(\{S_X \times S_Y \mid S_X \in \Sigma_X, S_Y \in \Sigma_Y\})$. The product of measurable spaces is then defined by $(X, \Sigma_X) \otimes (Y, \Sigma_Y) = (X \times Y, \Sigma_X \otimes \Sigma_Y)$.

Sum. Given measurable spaces (X, Σ_X) and (Y, Σ_Y) , we define a sum σ -algebra on the disjoint union $X + Y = \{(x, 0) \mid x \in X\} \cup \{(y, 1) \mid y \in Y\}$ by $\Sigma_X \oplus \Sigma_Y = \{S_X + S_Y \mid S_X \in \Sigma_X, S_Y \in \Sigma_Y\}$. The sum of measurable spaces is then defined by $(X, \Sigma_X) \oplus (Y, \Sigma_Y) = (X + Y, \Sigma_X \oplus \Sigma_Y)$.

Given measurable spaces X, Y and a measurable function $f: X \rightarrow Y$, define a new functor $\mathcal{L}: \mathbf{Meas} \rightarrow \mathbf{Meas}$ by $\mathcal{L}X = A \times X + 1$ along with its canonical σ -algebra $\Sigma_{\mathcal{L}X} = \mathcal{P}(A) \otimes \Sigma_X \oplus \mathcal{P}(1)$, and $\mathcal{L}f = id_A \times f + id_1$. Moreover, define $FX = [0, 1] \times [0, 1] \times X^A$ along with its σ -algebra $\mathcal{B}([0, 1]) \otimes \mathcal{B}([0, 1]) \otimes \bigotimes_{a \in A} \Sigma_X$ and $Ff = id_{[0, 1]} \times id_{[0, 1]} \times f^A$.

Integration. Let (X, Σ_X, m) be a measure space and $f: X \rightarrow \mathbb{R}$ be a measurable function. If $f(X) = \{\alpha_1, \dots, \alpha_n\}$ for some $\alpha_1, \dots, \alpha_n \in \mathbb{R}_+$, then f is called a *simple* function and its integral can be set as $\int_X f dm = \sum_{i=1}^n \alpha_i m(f^{-1}(\{\alpha_i\}))$. If f is non-negative, define $\int_X f dm = \sup \{ \int_X g dm \mid g \leq f, g \text{ simple} \} \in [0, \infty]$. Finally, for any $f: X \rightarrow \mathbb{R}$, decompose $f = f^+ - f^-$ where $f^+ \geq 0$ and $f^- \geq 0$. If their integrals are not both ∞ , define $\int_X f dm = \int_X f^+ dm - \int_X f^- dm$. If this is finite, we say that f is m -integrable. Furthermore, for any $S \in \Sigma_X$, the indicator function $\mathbf{1}_S$ is measurable and we define $\int_S f dm = \int_X \mathbf{1}_S f dm$.

Given a measurable function $g: X \rightarrow Y$ and measure $m: \Sigma_X \rightarrow \mathbb{R}_+$, the *pushforward measure* of m by g is $m \circ g^{-1}$. For any measurable $f: Y \rightarrow \mathbb{R}$, f is $m \circ g^{-1}$ -integrable iff $f \circ g$ is m -integrable and in this case, $\int_Y f d(m \circ g^{-1}) = \int_X (f \circ g) dm$. Each positive measurable function $X \rightarrow \mathbb{R}_+$ is the pointwise limit of an increasing sequence of simple functions. To prove some property for every positive measurable function, one can prove it for simple functions (or for indicator functions, if it is preserved by linear combinations) and show it is preserved by limits. Many such proofs use the monotone convergence theorem (see [14]), which states that if $(f_n)_{n \in \mathbb{N}}$ is an increasing sequence of positive functions with pointwise limit f , then f is measurable and $\int_X f dm = \lim \int_X f_n dm$.

The Giry Monad. The Giry monad [8] provides a link between probability theory and category theory. In $\mathbf{Meas}$, the Giry monad $(\mathbb{D}, \eta, \mu)$ is defined as follows. For any measurable space X , $\mathbb{D}X$ is the space of probability measures over (X, Σ_X) , and $\Sigma_{\mathbb{D}X}$ is the σ -algebra generated by the functions $e_S^X: \mathbb{D}X \rightarrow [0, 1]$ defined by $e_S^X(m) = m(S)$. For any measurable function $g: X \rightarrow Y$, $(\mathbb{D}g)(m) = m \circ g^{-1}$. The unit is defined by $\eta_X(x)(S) = \mathbf{1}_S(x)$ and the multiplication by $\mu_X(\Phi)(S) = \int_{\mathbb{D}X} e_S^X d\Phi$. Similarly, one defines the sub-Giry monad $(\mathbb{S}, \eta, \mu)$, with the only difference that $\mathbb{S}X$ is the space of sub-probability measures over (X, Σ_X) . There is a natural embedding of $\mathbb{D}$ in $\mathbb{S}$, denoted by $\iota: \mathbb{D} \Rightarrow \mathbb{S}$.

6.1 Trace Semantics via Determinisation

The aim of this section is to define trace semantics for continuous PTS, i.e., coalgebras of the form $\alpha: X \rightarrow \mathbb{D}(A \times X + 1)$ where X is a measurable space. We proceed in the same way as for discrete systems.

- (i) Transform α into a more convenient coalgebra $\tilde{\alpha}: X \rightarrow F\mathbb{S}X$.
- (ii) Determinise $\tilde{\alpha}$ into an F -coalgebra $\alpha^\sharp: \mathbb{S}X \rightarrow F\mathbb{S}X$.
- (iii) Factorise the final morphism $\varphi_{\alpha^\sharp} = \varphi_\Pi \circ [-]$, then precompose with η_X .

The following diagram sums up the construction. Here $\Sigma_{([0,1] \times [0,1])^{A^*}}$ is the Σ -algebra generated by the functions $L \mapsto L(w)$.

$$\begin{array}{ccccccc}
 X & \xrightarrow{\eta_X} & \mathbb{S}X & \xrightarrow{\varphi} & (\mathbb{S}A^\infty)^\infty & \xrightarrow{\omega} & ([0,1] \times [0,1])^{A^*} \\
 \downarrow \alpha & \searrow \tilde{\alpha} & \downarrow \alpha^\sharp & \downarrow [-] & \downarrow \Pi & & \downarrow \omega \\
 \mathbb{D}(A \times X + 1) & & \mathbb{S}(A \times X + 1) & \xrightarrow{\epsilon_X} & [0,1] \times [0,1] \times (\mathbb{S}X)^A & \xrightarrow{\text{id} \times (\varphi)^A} & [0,1] \times [0,1] \times (([0,1] \times [0,1])^{A^*})^A \\
 \downarrow \iota_{A \times X + 1} & & & & & & \\
 \mathbb{S}(A \times X + 1) & & & & & &
 \end{array}$$

(i) Translation: from α to $\tilde{\alpha}$

Proposition 6.1. *For any measurable space X , the function $\epsilon_X: \mathbb{S}\mathbb{L}X \rightarrow F\mathbb{S}X$ defined by $\epsilon_X(m) = \langle m(\mathbb{L}X), m(1), a \mapsto [S \mapsto m(\{a\} \times S)] \rangle$ is measurable. Moreover, $\epsilon: \mathbb{S}\mathbb{L} \Rightarrow F\mathbb{S}$ is a natural transformation.*

Now take $\tilde{\alpha} = \epsilon_X \circ \iota_{\mathbb{L}X} \circ \alpha: X \rightarrow [0,1] \times [0,1] \times (\mathbb{S}X)^A$. Explicitly:

$$\tilde{\alpha}(x) = \langle \underbrace{\alpha(x)(\mathbb{L}X)}_1, \alpha(x)(1), a \mapsto [S \mapsto \alpha(x)(\{a\} \times S)] \rangle$$

We decompose it as a pairing $\tilde{\alpha} = \langle \tilde{\alpha}_\oplus, \tilde{\alpha}_*, a \mapsto t_a \rangle$.

(ii) Determinisation. We recall some basic observations in abstract determinisation [12, 16]. By *distributive law* here we mean the standard notion of distributive law of monad over functor (called EM-law in [12]).

Lemma 6.2. *Let $\mathcal{C}$ be a category, $F: \mathcal{C} \rightarrow \mathcal{C}$ be an endofunctor and (T, η, μ) be a monad on $\mathcal{C}$. Let $f: X \rightarrow TFX$ be a TF -coalgebra and $h: TFTX \rightarrow FTX$ be an Eilenberg-Moore T -algebra. Then there exists a unique T -algebra morphism $f^\sharp: (TX, \mu_X) \rightarrow (FTX, h)$ such that $f = f^\sharp \circ \eta_X$.*

Lemma 6.3. *With the same notations as for Lemma 6.2, and given a distributive law $\lambda: TF \Rightarrow FT$, then $h = F\mu_X \circ \lambda_{TX}: TFTX \rightarrow FTX$ is an Eilenberg-Moore T -algebra.*

The next step is to define a distributive law $\lambda: \mathbb{S}F \Rightarrow F\mathbb{S}$ in order to apply Lemmas 6.2 and 6.3. In the following we write $id_{FX} = \langle \pi_X^\oplus, \pi_X^*, a \mapsto \pi_X^a \rangle$. Note that $\pi^\epsilon: F \Rightarrow [0, 1]$ (for $\epsilon \in \{*, \oplus\}$) and $\pi^a: F \Rightarrow Id_{\mathbf{C}}$ (for $a \in A$) are natural transformations.

Lemma 6.4. *Let $g: \mathbb{S}([0, 1]) \rightarrow [0, 1]$ be defined by $g(m) = \int_{[0, 1]} id_{[0, 1]} dm$. Then g is measurable and an Eilenberg-Moore $\mathbb{S}$ -algebra.*

For any object X of **Meas**, define $\lambda_X: \mathbb{S}FX \rightarrow F\mathbb{S}X$ by

$$\lambda_X = \langle g \circ \mathbb{S}\pi_X^\oplus, g \circ \mathbb{S}\pi_X^*, a \mapsto \mathbb{S}\pi_X^a \rangle$$

This is a measurable function because each component is measurable.

Proposition 6.5. *$\lambda: \mathbb{S}F \Rightarrow F\mathbb{S}$ is a distributive law.*

Let us compute the value of our resulting determinisation. Given $\tilde{\alpha}: X \rightarrow F\mathbb{S}X$, take $h = F\mu_X \circ \lambda_{\mathbb{S}X}$ (Lemma 6.3) and $\alpha^\sharp = h \circ \mathbb{S}\tilde{\alpha}$ (Lemma 6.2). We get

$$\begin{aligned} \alpha^\sharp &= h \circ \mathbb{S}\tilde{\alpha} \\ &= F\mu_X \circ \lambda_{\mathbb{S}X} \circ \mathbb{S}\tilde{\alpha} \\ &= F\mu_X \circ \langle g \circ \mathbb{S}(\pi_{\mathbb{S}X}^\oplus \circ \tilde{\alpha}), g \circ \mathbb{S}(\pi_{\mathbb{S}X}^* \circ \tilde{\alpha}), a \mapsto \mathbb{S}(\pi_{\mathbb{S}X}^a \circ \tilde{\alpha}) \rangle \\ &= \langle g \circ \mathbb{S}\tilde{\alpha}_\oplus, g \circ \mathbb{S}\tilde{\alpha}_*, a \mapsto \mu_X \circ St_a \rangle \end{aligned}$$

Let $m \in \mathbb{S}X$. This more explicit expression shows that the coalgebra that arises from the determinisation is natural in the sense that the components of $\alpha^\sharp$ are basically obtained by integrating the information provided by α .

$$\begin{aligned} \alpha^\sharp(m) &= \left\langle \int_X \tilde{\alpha}_\oplus dm, \int_X \tilde{\alpha}_* dm, a \mapsto \left[S \mapsto \int_X t_a(-)(S) dm \right] \right\rangle \\ &= \left\langle \int_X \alpha(-)(\mathcal{L}X) dm, \int_X \alpha(-)(1) dm, a \mapsto \left[S \mapsto \int_X \alpha(-)(\{a\} \times S) dm \right] \right\rangle \end{aligned}$$

(iii) **Final coalgebra.** This heavy determinisation part gives us an F -coalgebra $\alpha^\sharp$. There exists a final object in $\text{CoAlg}(F)$:

Proposition 6.6. *Let $\Omega = ([0, 1] \times [0, 1])^{A^*}$ and Σ_Ω be the smallest σ -algebra that makes the functions $e_w: \Omega \rightarrow [0, 1] \times [0, 1]$ defined by $e_w(L) = (L(w), \text{measurable for every } w \in A^*)$. Let $\omega: \Omega \rightarrow F\Omega$ be defined by $\omega(L) = \langle L(\varepsilon), a \mapsto L_a \rangle$. Then (Ω, ω) is a final F -coalgebra.*

Thus for any F -coalgebra β the final morphism towards ω , denoted φ_β , gives a canonical notion of semantics. What we want is something slightly more specific that takes into account the way $\alpha^\sharp$ was built to produce a probability measure in $\mathbb{S}A^\infty$. This is obtained via the coalgebra $\Pi: \mathbb{S}A^\infty \rightarrow F\mathbb{S}A^\infty$, built as follows.

Proposition 6.7. *Let $\pi: A^\infty \rightarrow \mathcal{L}A^\infty$ be defined by $\pi(\varepsilon) = *$ and $\pi(aw) = (a, w)$. This is a final $\mathcal{L}$ -coalgebra.*

Let $\Pi = \mathbf{e}_{A^\infty} \circ \mathbb{S}\pi$. One can check that with this definition, $\Pi : \mathbb{S}A^\infty \rightarrow F\mathbb{S}A^\infty$ has the same expression as the $\Pi : \mathcal{M}(A^\infty) \rightarrow F\mathcal{M}(A^\infty)$ of Sect. 4:

$$\begin{aligned}\Pi(m) &= \langle m(\pi^{-1}(\mathfrak{L}A^\infty)), m(\pi^{-1}(1)), a \mapsto [S \mapsto m(\pi^{-1}(\{a\} \times S))] \rangle \\ &= \langle m(A^\infty), m(\varepsilon), a \mapsto m_a \rangle\end{aligned}$$

The aim is now to factorise the semantics obtained via ω into semantics obtained via Π . The following result is a kind of completeness property for this operation.

Lemma 6.8. *The final morphism φ_Π from Π to ω is injective.*

Proof. For any $m, m' \in \mathbb{S}A^\infty$, in order to have $m = m'$, it is sufficient that $m|_{S_\infty} = m'|_{S_\infty}$ (see Theorem 3.5). By induction on w , we prove that for $m, m' \in \mathbb{S}A^\infty$ such that $\varphi_\Pi(m) = \varphi_\Pi(m')$, then $\langle m(wA^\infty), m(w) \rangle = \langle m'(wA^\infty), m'(w) \rangle$. First, $\langle m(\varepsilon A^\infty), m(\varepsilon) \rangle = \varphi_\Pi(m)(\varepsilon) = \varphi_\Pi(m')(\varepsilon) = \langle m'(\varepsilon A^\infty), m'(\varepsilon) \rangle$. Note that $\varphi_\Pi(m) = \varphi_\Pi(m')$ implies $\varphi_\Pi(m_a)(w) = \varphi_\Pi(m)(aw) = \varphi_\Pi(m')(aw) = \varphi_\Pi(m'_a)(w)$ so that $\varphi_\Pi(m_a) = \varphi_\Pi(m'_a)$. Use the induction hypothesis to obtain that $\langle m(awA^\infty), m(aw) \rangle = \langle m_a(wA^\infty), m_a(w) \rangle = \langle m'_a(wA^\infty), m'_a(w) \rangle = \langle m'(awA^\infty), m'(aw) \rangle$. This achieves the induction, so m and m' coincide on S_∞ , hence $m = m'$. $\square$

The following proposition states precisely in which cases the factorisation can be done. This is a variant of Theorem 3.5 in which we really see that our system is making one step. This version is stronger than Proposition 4.5, because it also proves that the involved functions are measurable.

Theorem 6.9. *Let $\beta = \langle \beta_\oplus, \beta_*, a \mapsto \tau_a \rangle : Y \rightarrow FY$ be an F -coalgebra. The two following conditions are equivalent:*

- (i) *There exists an F -coalgebra morphism $[-]$ from β to Π .*
- (ii) *The equation $\beta_\oplus = \beta_* + \sum_{a \in A} \beta_\oplus \circ \tau_a$ holds.*

In this case, this morphism is unique.

For convenience we denote $e_S^{A^\infty} \circ [-]$ by $[-](S)$, and $\phi_a \circ [-]$ by $[-]_a$, where the measure derivative function $\phi_a : m \mapsto m_a$ is measurable as a component of Π .

Proof. (i) $\Rightarrow$ (ii). Suppose $[-]$ is a coalgebra morphism from β to Π . Commutation of the diagram yields $\langle \beta_\oplus, \beta_*, a \mapsto [-] \circ \tau_a \rangle = \langle [-](A^\infty), [-](\varepsilon), a \mapsto [-]_a \rangle$. Let $y \in Y$. Because $[y]$ is a measure, $\beta_\oplus(y) = [y](\varepsilon A^\infty) = [y](\varepsilon) + \sum_{a \in A} [y](aA^\infty)$. Thus $\beta_\oplus(y) = \beta_*(y) + \sum_{a \in A} [\tau_a(y)](A^\infty) = \beta_*(y) + \sum_{a \in A} (\beta_\oplus \circ \tau_a)(y)$.

Uniqueness. If $[-]'$ is another such morphism, we have $[-](A^\infty) = [-]'(A^\infty)$, $[-](\varepsilon) = [-]'(\varepsilon)$ and for any $a \in A$, $[-] \circ \tau_a = [-]_a$ and $[-]' \circ \tau_a = [-]'_a$. An immediate induction yields $[-]|_{S_\infty} = [-]'|_{S_\infty}$, thus $[-] = [-]'$ by Theorem 3.5.

(ii) $\Rightarrow$ (i) Assume that (ii) holds. Let us define $[-]$ on S_∞ by induction:

$$\begin{aligned}[y]|_{S_\infty}(\varepsilon A^\infty) &= \beta_\oplus(y) & [y]|_{S_\infty}(\varepsilon) &= \beta_*(y) \\ [y]|_{S_\infty}(awA^\infty) &= [\tau_a(y)]|_{S_\infty}(wA^\infty) & [y]|_{S_\infty}(aw) &= [\tau_a(y)]|_{S_\infty}(w)\end{aligned}$$

We must prove that it can be extended to a measure, using Theorem 3.5. First, note that $[y]_{|S_\infty}(\varepsilon A^\infty) = \beta_\oplus(y) = \beta_*(y) + \sum_{a \in A} (\beta_\oplus \circ \tau_a)(y) = [y]_{|S_\infty}(\varepsilon) + \sum_{a \in A} [y]_{|S_\infty}(aA^\infty)$. If it is known that for all $y \in Y$, $[y]_{|S_\infty}(wA^\infty) = [y]_{|S_\infty}(w) + \sum_{a \in A} [y]_{|S_\infty}(waA^\infty)$ then for any $b \in A$ we obtain the equation $[y]_{|S_\infty}(bwA^\infty) = [\tau_b(y)]_{|S_\infty}(wA^\infty) = [\tau_b(y)]_{|S_\infty}(w) + \sum_{a \in A} [\tau_b(y)]_{|S_\infty}(waA^\infty) = [y]_{|S_\infty}(bw) + \sum_{a \in A} [y]_{|S_\infty}(bwaA^\infty)$. This proves the (ii) of Theorem 3.5. We denote by $[-]$ the extension of $[-]_{|S_\infty}$. We postpone the proof of the measurability of $[-]$; what is left is the commutation of the coalgebra diagram. The first line of the definition of $[-]_{|S_\infty}$ gives directly that $\beta_\oplus = [-](A^\infty)$ and $\beta_* = [-](\varepsilon)$. Let $a \in A$. For any $y \in Y$, according to the second line of the definition of $[-]_{|S_\infty}$, the measures $[\tau_a(y)]$ and $[y]_a$ coincide on S_∞ , hence are equal according to Theorem 3.5, so $[-] \circ \tau_a = [-]_a$. This achieves the proof that the diagram commutes.

Measurability. It is not immediate to see why $[-]: Y \rightarrow \mathbb{S}A^\infty$ is a measurable function. What has to be shown is that for any $S \in \Sigma_{A^\infty}$, $[-](S)$ is measurable. This is true when $S \in S_\infty$. Indeed, $[-](\emptyset)$ is the zero function, which is measurable. For the rest we proceed by induction. Obviously $[-](\varepsilon A^\infty) = \beta_\oplus$ and $[-](\varepsilon) = \beta_*$ are measurable because β is. Furthermore, $[-](awA^\infty) = [-]_a(wA^\infty) = [-](wA^\infty) \circ \tau_a$ and $[-](aw) = [-]_a(w) = [-](w) \circ \tau_a$ are measurable by induction hypothesis and composition.

We need to introduce a widely known theorem of measure theory, namely the $\pi - \lambda$ theorem (see [1], Lemma 4.11). Let Z be a set. A set $P \subseteq \mathcal{P}(Z)$ is a π -system if it is non-empty and closed under finite intersections. A set $D \subseteq \mathcal{P}(Z)$ is a λ -system if it contains Z and is closed under difference (if $A, B \in D$ and $A \subseteq B$ then $B \setminus A \in D$) and countable increasing union. The $\pi - \lambda$ theorem states that given P a π -system, D a λ -system such that $P \subseteq D$, then $\sigma_Z(P) \subseteq D$.

Take $Z = A^\infty$, $P = S_\infty$ and $D = \{S \in \Sigma_{A^\infty} \mid [-](S) \text{ is measurable}\}$. It is easy to see that S_∞ is a π -system. Moreover, D is a λ -system. Indeed, $A^\infty \in D$ (see above), if $(S_n)_{n \in \mathbb{N}}$ is an increasing sequence of sets in D , then $[-](S_1 \setminus S_0) = [-](S_1) - [-](S_0)$ is measurable as a difference of measurable functions and $[-](\bigcup_{n \in \mathbb{N}} S_n) = \lim_{n \rightarrow \infty} [-](S_n)$ is measurable as a pointwise limit of measurable functions. Finally, given the preceding paragraph, we have $S_\infty \subseteq D$. The $\pi - \lambda$ theorem therefore yields $\Sigma_{A^\infty} \subseteq D$. Thus $[-]$ is measurable. $\square$

An interpretation of the last proposition is that, in the subcategory of all F -coalgebras that satisfy the equation (ii), the final object is Π . If Theorem 6.9 holds, then note that $\varphi_\Pi \circ [-]$ is a coalgebra morphism from β into the final coalgebra ω . Hence by finality $\varphi_\Pi \circ [-] = \varphi_\beta$. Along with Lemma 6.8, this yields the following proposition, which is exactly the same as in Sect. 2.

Proposition 6.10. *Let $\beta: Y \rightarrow FY$ be an F -coalgebra for which Theorem 6.9 holds. Then for any $y, z \in Y$, $[y] = [z]$ iff $\varphi_\beta(y) = \varphi_\beta(z)$.*

Back to $\alpha: X \rightarrow \mathbb{D}\mathcal{L}X$ we check that Theorem 6.9 holds for $\alpha^\sharp = \langle \alpha_\oplus^\sharp, \alpha_*^\sharp, a \mapsto \tau_a \rangle$. Note that because $\alpha(-)(\mathcal{L}X) = 1$, we have for all $m \in \mathbb{S}X$ that $m(X) =$

$\int_X 1 dm = \int_X \alpha(-)(\mathfrak{L}X) dm = \alpha_{\oplus}^{\sharp}(m)$. This justifies the last equality:

$$\begin{aligned} \alpha_{\oplus}^{\sharp}(m) &= \int_X \alpha(-)(\mathfrak{L}X) dm = \int_X \left(\alpha(-)(1) + \sum_{a \in A} \alpha(-)(\{a\} \times X) \right) dm \\ &= \int_X \alpha(-)(1) dm + \sum_{a \in A} \int_X \alpha(-)(\{a\} \times X) dm \\ &= \alpha_{*}^{\sharp}(m) + \sum_{a \in A} \tau_a(m)(X) = \alpha_{*}^{\sharp}(m) + \sum_{a \in A} (\alpha_{\oplus}^{\sharp} \circ \tau_a)(m) \end{aligned}$$

Conclusion. Any $\alpha: X \rightarrow \mathbb{D}\mathfrak{L}X$ can be given a canonical trace semantics via a determinisation process. This is a function $[-]: \mathbb{S}X \rightarrow \mathbb{S}A^{\infty}$.

6.2 Correctness of the Resulting Trace Semantics

In [13], given a PTS $\alpha: X \rightarrow \mathbb{D}\mathfrak{L}X$, the trace semantics $\llbracket - \rrbracket: X \rightarrow \mathbb{D}A^{\infty}$ (denoted by $\mathbf{tr}$ in [13]) is defined by

$$\begin{aligned} \llbracket x \rrbracket(\varepsilon A^{\infty}) &= \alpha(x)(\mathfrak{L}X) \quad (= 1) & \llbracket x \rrbracket(\varepsilon) &= \alpha(x)(1) \\ \llbracket x \rrbracket(awA^{\infty}) &= \int_X \llbracket - \rrbracket(wA^{\infty}) dt_a(x) & \llbracket x \rrbracket(aw) &= \int_X \llbracket - \rrbracket(w) dt_a(x) \end{aligned}$$

We will hereby prove that this semantics fits with ours, in the sense that the following diagram commutes.

$$\begin{array}{ccc} X & \xrightarrow{\llbracket - \rrbracket} & \mathbb{D}A^{\infty} \\ \eta_X \downarrow & & \downarrow \iota_{A^{\infty}} \\ \mathbb{S}X & \xrightarrow{[-]} & \mathbb{S}A^{\infty} \end{array} \quad (4)$$

Lemma 6.11. *For any $m \in \mathbb{S}X$ and any $S \in S_{\infty}$, $[m](S) = \int_X ([-] \circ \eta_X)(S) dm$.*

Proof. In this proof, $\int_X f dm$ may be denoted by $\int_{x \in X} f(x) m(dx)$. One can show using the monotone convergence theorem that for any measurable function $f: X \rightarrow [0, 1]$,

$$\int_X f d\tau_a(m) = \int_{x \in X} \left(\int_X f dt_a(x) \right) m(dx)$$

Note further that $[\eta_X(x)](\varepsilon A^{\infty}) = (\alpha_{\oplus}^{\sharp} \circ \eta_X)(x) = \tilde{\alpha}_{\oplus}(x) = \alpha(x)(\mathfrak{L}X)$ and in the same way $[\eta_X(x)](\varepsilon) = \alpha(x)(1)$. Now let us prove the lemma by induction, for all $m \in \mathbb{S}X$. First

$$\begin{aligned} [m](\varepsilon A^{\infty}) &= \alpha_{\oplus}^{\sharp}(m) = \int_X \alpha(-)(\mathfrak{L}X) dm = \int_X ([-] \circ \eta_X)(\varepsilon A^{\infty}) dm \\ [m](\varepsilon) &= \alpha_{*}^{\sharp}(m) = \int_X \alpha(-)(1) dm = \int_X ([-] \circ \eta_X)(\varepsilon) dm \end{aligned}$$

Assume the result is true for wA^∞ and w . Take $\diamond \in \{\{\varepsilon\}, A^\infty\}$.

$$\begin{aligned}
[m](aw\diamond) &= [\tau_a(m)](w\diamond) = \int_X ([-] \circ \eta_X)(w\diamond) d\tau_a(m) && \text{(induction hypothesis)} \\
&= \int_{x \in X} \left(\int_X ([-] \circ \eta_X)(w\diamond) dt_a(x) \right) m(dx) && \text{(preliminary remark)} \\
&= \int_{x \in X} [\tau_a(\eta_X(x))](w\diamond) m(dx) && \text{(definition of } \tau_a) \\
&= \int_X ([-] \circ \eta_X)(aw\diamond) dm
\end{aligned}$$

□

Using this last lemma and that $\tau_a \circ \eta_X = t_a$, we have for any $x \in X$:

$$\begin{aligned}
[\eta_X(x)](\varepsilon A^\infty) &= \alpha(x)(\mathfrak{L}X) \\
[\eta_X(x)](\varepsilon) &= \alpha(x)(1) \\
[\eta_X(x)](awA^\infty) &= [(\tau_a \circ \eta_X)(x)](wA^\infty) \\
&= [t_a(x)](wA^\infty) = \int_X [\eta_X(-)](wA^\infty) dt_a(x) \\
[\eta_X(x)](aw) &= [(\tau_a \circ \eta_X)(x)](w) = [t_a(x)](w) = \int_X [\eta_X(-)](w) dt_a(x)
\end{aligned}$$

Thus, for any $x \in X$, $([-] \circ \eta_X)(x)$ and $(\iota_{A^\infty} \circ \llbracket - \rrbracket)(x)$ are measures in $\mathbb{S}A^\infty$ that coincide on S_∞ . Because of Theorem 3.5, they are equal. Consequently, the trace semantics we get via determinisation and Eilenberg-Moore algebras is the same as the Kleisli trace semantics of [13].

Theorem 6.12. *The diagram (4) commutes, i.e., the maps $\iota_{A^\infty} \circ \llbracket - \rrbracket$ and $[-] \circ \eta_X$ coincide.*

Finally, note that, in the event that $\alpha: X \rightarrow \mathbb{D}\mathfrak{L}X$ can be seen as a discrete system, i.e., for all $x \in X$, $\alpha(x)$ is a convex countable sum of Dirac distributions, then the general semantics coincide with those obtained in Sect. 2.

7 Related Work

The (in)finite trace semantics of PTS discussed in this paper was presented coalgebraically in [13], through the Kleisli category of the (sub-)Giry monad. By using a determinisation construction, we obtain the same trace semantics, in a fundamentally different way. This determinisation construction is precisely what allows us to use bisimulations (up-to) to prove equivalence. Further, our determinisation construction can be presented separately for the discrete and continuous cases (the discrete case is much more basic), whereas in the Kleisli setting only the general continuous case can be presented (since discrete systems generate a probability measure). Other coalgebraic approaches to infinite traces (based on fixed points, e.g., [7, 19]) do not use determinisation.

Our determinisation construction for (in)finite traces is strongly inspired by the one for finite traces in [12, 17]. As explained in Sect. 4, the main technical

difference is that the total probability mass of states in the determinised Moore automaton becomes observable, and that this yields a probability measure over sets of traces rather than a (sub)probability distribution over individual traces.

The above-mentioned equivalence between the determinisation and ‘Kleisli’ trace semantics for finite traces is a motivating example for the general comparison between coalgebraic determinisation and Kleisli traces in [12]. However, we do not know if those results can be applied here for at least one reason: the correspondence stated in [12] uses only one monad for both constructions, using, in case of finite traces, an extension natural transformation of the form $\epsilon: \mathbb{S}\mathbb{L} \Rightarrow F\mathbb{S}$ (actually, the discrete version). However, in our construction, we have to move from probability measures in the definition of PTSs (modeled by $\mathbb{D}$) to sub-probability measures in the determinised Moore automaton (modeled by $\mathbb{S}$). In contrast to the case of finite traces, we can not simply replace $\mathbb{D}$ by $\mathbb{S}$ in the definition of PTS, since the sums-to-1 condition is required for the condition (ii) of Theorem 6.9. One might try to nevertheless use only $\mathbb{S}$ as the monad, focusing on PTSs (involving $\mathbb{S}$) that satisfy the sums-to-one condition. But it is currently unclear to us how such a subclass fits into the framework of [12]; moreover, the Kleisli semantics for PTSs based on $\mathbb{S}$ is finite traces [13, Theorem 3.33]. Another idea is to use the isomorphism $\mathbb{D}(A \times X + 1) \simeq \mathbb{S}(A \times X)$, (via the map $m \mapsto m|_{\Sigma_A \otimes \Sigma_X}$) but this does not seem to solve the issue: the Kleisli semantics of PTS of the form $X \rightarrow \mathbb{S}(A \times X)$ is trivial [13, Theorem 3.33]. We leave a suitable extension of the abstract framework [12] for future work.

For the algorithm presented in Sect. 5, we embed convex combinations (in the transition structure of PTS) into vector spaces, in order to use a more general contextual closure, w.r.t. arbitrary linear combinations rather than only convex combinations. This guarantees termination of the algorithm based on bisimulation up to congruence. We do not know whether this move is really necessary: perhaps the contextual closure w.r.t. only convex combinations suffices. The recent [5] might be of use in answering this question.

This work was done primarily from a coalgebraic point of view. Actually, as pointed out by one of the reviewers, the determinization of a PTS involves to a standard construction in the theory of Markov chains and stochastic processes: the passage from a kernel to a stochastic operator. This perspective could be investigated further. Notably, one motivation for trying to do so is to study how the results of Sect. 5 could extend to (discrete approximations) of the measurable PTSs of Sect. 6.

Acknowledgments. We are grateful to Filippo Bonchi, Paul Levy, Damien Pous, Jan Rutten, Ana Sokolova and the anonymous reviewers for comments and suggestions.

References

1. Aliprantis, C.D., Border, K.C.: Infinite Dimensional Analysis: A Hitchhiker’s Guide. Springer, Berlin (2006). <https://doi.org/10.1007/3-540-29587-9>
2. Bonchi, F., König, B., Küpper, S.: Up-to techniques for weighted systems. In: Legay, A., Margaria, T. (eds.) TACAS 2017. LNCS, vol. 10205, pp. 535–552. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54577-5_31

3. Bonchi, F., Petrisan, D., Pous, D., Rot, J.: A general account of coinduction up-to. *Acta Inf.* **54**(2), 127–190 (2017). <https://doi.org/10.1007/s00236-016-0271-4>
4. Bonchi, F., Pous, D.: Checking NFA equivalence with bisimulations up to congruence. In: *Principle of Programming Languages (POPL)*, Roma, Italy, pp. 457–468. ACM (2013). <https://doi.org/10.1145/2429069.2429124>. 16 p
5. Bonchi, F., Silva, A., Sokolova, A.: The power of convex algebras. In: Meyer, R., Nestmann, U. (eds.) *28th International Conference on Concurrency Theory (CONCUR 2017)*. *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 85, pp. 23:1–23:18. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2017). <https://doi.org/10.4230/LIPIcs.CONCUR.2017.23>
6. Boreale, M.: Weighted bisimulation in linear algebraic form. In: Bravetti, M., Zavattaro, G. (eds.) *CONCUR 2009*. *LNCS*, vol. 5710, pp. 163–177. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-04081-8_12
7. Cirstea, C.: From branching to linear time, coalgebraically. *Fundam. Inform.* **150**(3–4), 379–406 (2017). <https://doi.org/10.3233/FI-2017-1474>
8. Giry, M.: A categorical approach to probability theory. *Lect. Notes Math.* **915**, 68–85 (1982)
9. Goy, A.: Trace semantics via determinization for probabilistic transition systems. Master internship report, Radboud University (2017). [arxiv:1802.09084](https://arxiv.org/abs/1802.09084)
10. Hasuo, I., Jacobs, B., Sokolova, A.: Generic trace semantics via coinduction. *Logical Methods Comput. Sci.* **3**(4) (2007). [https://doi.org/10.2168/LMCS-3\(4:11\)](https://doi.org/10.2168/LMCS-3(4:11))
11. Jacobs, B.: *Introduction to Coalgebra: Towards Mathematics of States and Observation*, vol. 59. Cambridge University Press, Cambridge (2016). <https://doi.org/10.1017/CBO9781316823187>. Cambridge Tracts in Theoretical Computer Science
12. Jacobs, B., Silva, A., Sokolova, A.: Trace semantics via determinization. *J. Comput. Syst. Sci.* **81**(5), 859–879 (2015). <https://doi.org/10.1016/j.jcss.2014.12.005>. 11th International Workshop on Coalgebraic Methods in Computer Science, CMCS 2012 (Selected Papers)
13. Kerstan, H., König, B.: Coalgebraic trace semantics for continuous probabilistic transition systems. *Logical Methods Comput. Sci.* **9**(4) (2013). [https://doi.org/10.2168/LMCS-9\(4:16\)2013](https://doi.org/10.2168/LMCS-9(4:16)2013)
14. Klenke, A.: *Probability Theory: A Comprehensive Course*. Springer, London (2013). <https://doi.org/10.1007/978-1-4471-5361-0>
15. Rutten, J.: Universal coalgebra: a theory of systems. *TCS* **249**(1), 3–80 (2000)
16. Silva, A., Bonchi, F., Bonsangue, M.M., Rutten, J.J.M.M.: Generalizing determinization from automata to coalgebras. *Log. Methods Comput. Sci.* **9**(1) (2013). [https://doi.org/10.2168/LMCS-9\(1:9\)2013](https://doi.org/10.2168/LMCS-9(1:9)2013)
17. Silva, A., Sokolova, A.: Sound and complete axiomatization of trace semantics for probabilistic systems. *Electr. Notes Theor. Comput. Sci.* **276**, 291–311 (2011). <https://doi.org/10.1016/j.entcs.2011.09.027>. Twenty-seventh Conference on the Mathematical Foundations of Programming Semantics (MFPS XXVII)
18. Sokolova, A.: Probabilistic systems coalgebraically: a survey. *Theor. Comput. Sci.* **412**(38), 5095–5110 (2011). <https://doi.org/10.1016/j.tcs.2011.05.008>
19. Urabe, N., Hasuo, I.: Coalgebraic infinite traces and Kleisli simulations. In: Moss, L.S., Sobocinski, P. (eds.) *CALCO 2015*. *LIPIcs*, vol. 35, pp. 320–335. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2015). <https://doi.org/10.4230/LIPIcs.CALCO.2015.320>