Math. Struct. in Comp. Science (2011), *vol.* 21, *pp.* 175–234. © Cambridge University Press 2011
doi:10.1017/S0960129510000526

Coalgebraic logic over general measurable spaces – a survey

ERNST-ERICH DOBERKAT and CHRISTOPH SCHUBERT

Chair for Software Technology, Technische Universität Dortmund
Email: {ernst-erich.doberkat;christoph.schubert}@udo.edu

Received 31 March 2010; revised 15 September 2010

In this survey we discuss the generalisation of stochastic Kripke models for general modal logics through predicate liftings for functors over general measurable spaces. We derive results on expressivity and show that selection arguments allow us to incorporate the discussion of bisimilarity, provided the underlying spaces are assumed to be standard Borel.

Contents

1 Introduction 176
2 Modal logics 178
2.1 Definitions 179
2.2 Kripke models as coalgebras 180
2.3 Bisimilar Kripke models 181
3 Measure theoretic and other preliminaries 182
3.1 Measurable spaces 183
3.2 Measurable morphisms 183
3.3 Subprobability functor 185
3.4 A small dose of topology 187
3.5 Coalgebra 191
4 Coalgebraic logic using predicate liftings 192
4.1 Motivation 193
4.2 Examples of predicate liftings 195
4.3 The problem of expressivity 196
4.4 Some general results for coalgebraic logic 198
4.5 Expressive coalgebraic logic over **Set** 199
5 Logical equivalence versus behavioural equivalence 199
5.1 Generating and separating sets of predicate liftings 200
5.2 The congruence theorem 201
5.3 Expressivity 203
5.4 Closure properties of predicate liftings 205
5.5 Terminal coalgebras 210
5.6 Examples 211
6 Bisimilarity 215
6.1 Selections 215

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

6.2 Bisimilarity 217
6.3 Semi-pullbacks exist 218
6.4 Congruences and bisimulations 221
6.5 Behavioural equivalence 226
7 Further research 230
References 232

1. Introduction

As motivation for the problems we will consider here, consider the following very simple logic with formulas given by

$$\phi ::= \top \mid \phi_1 \wedge \phi_2 \mid \langle a \rangle_q \phi.$$

Here $a \in \mathsf{Act}$ is an action, and $q \in \mathbb{Q} \cap [0,1]$ is a rational number. The informal probabilistic interpretation of this logic reads as follows: formula $\langle a \rangle_q \phi$ holds in a state s if the probability of entering a state s' in which ϕ holds after action a is bounded below by q. Now assume that a Kripke model $(S, (k_a)_{a \in \mathsf{Act}})$ for interpreting the logic is given. Here S is a state space, and k_a is a transition probability, $k_a(s)(A)$ being the probability that the system's new state after action a in state s is a member of $A \subseteq S$; we may, and do, assume that we are dealing with subprobabilities, so $k_a(s)(S) \leqslant 1$ for all $a \in \mathsf{Act}, s \in S$. Let $[\![\phi]\!]$ be the set of all states in which formula ϕ holds. Then the informal interpretation from above translates into

$$[\![\langle a \rangle_q \phi]\!] = \{\, s \in S \mid k_a(s)([\![\phi]\!]) \geqslant q \,\}. \tag{1}$$

This way of interpreting modal logics was first proposed in Larsen and Skou (1991) for discrete state spaces S, and was later generalised by adopting continuous state spaces. The seminal paper Desharnais *et al.* (2002) assumes that the state space is an analytic space, that is, the continuous image of a Polish space, and this assumption was also made in Doberkat (2006). The reason for this topological assumption is that it makes semi-pullbacks for stochastic Kripke models available, which in turn is helpful for investigating questions of expressivity. We will touch on this point in Section 6. When investigating other logics, in particular logics for model checking (Baier *et al.* 2003), the methods developed for the Hennessy–Milner logic from above can be taken over nearly verbatim (Doberkat 2007a). This is quite surprising, since the logics appear fairly different in syntax, and in scope.

Looking at equation (1) again, it can be rewritten to make the main actors more visible in the following way:

$$\begin{aligned}
[\![\langle a \rangle_q \phi]\!] &= \{\, s \in S \mid k_a(s)([\![\phi]\!]) \geqslant q \,\} \\
&= k_a^{-1}\big[\{\, \mu \in \mathfrak{S}(S) \mid \mu([\![\phi]\!]) \geqslant q \,\}\big] \\
&= k_a^{-1}\big[\varpi_{q,S}([\![\phi]\!])\big] \\
&= \big(k_a^{-1} \cdot \varpi_{q,S}\big)([\![\phi]\!]).
\end{aligned}$$

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Here

$$\varpi_{q,S}(A) := \{\, \mu \in \mathfrak{S}(S) \mid \mu(A) \geqslant q \,\}$$

with $\mathfrak{S}(S)$ as the set of all subprobabilities over S. In the last line, k_a^{-1} denotes the inverse-image map of $k_a : S \to \mathfrak{S}(S)$ – it assigns a subset of $\mathfrak{S}(S)$ to a subset of S. Similarly, $(\varpi_{q,S})_S$ is viewed as a natural transformation from the Borel sets of the state space S to the Borel sets of the subprobabilities on S. Observe that we can consider k_a to be the dynamics of a coalgebra for the subprobability functor $\mathfrak{S}$; see Section 3.5. Hence, we interpret the modal formula through $k_a^{-1} \cdot \varpi_{q,S}$, *viz*. the inverse-image map of the dynamics of a coalgebra underlying the Kripke model in combination with a natural transformation in which the functor governing the coalgebra plays a dominant role.

This observation leads to coalgebraic logic in which a set of natural transformations involving a functor, which in turn governs a coalgebra's dynamics, helps in interpreting the logic, the formulas of which are somewhat geared towards these transformations. In contrast to Moss (1999), Cîrstea (2004), Pattinson (2004), Schröder and Pattinson (2007) and Schröder (2008), to name just a few, we do not work with a functor based on the category of sets. We take the category of measurable spaces with measurable maps as our base category, and consider an endofunctor there.

For specialised cases, *viz*. if we work in a Polish or analytic space, expressivity for this kind of coalgebraic logic has been investigated in Doberkat and Schubert (2009) and Doberkat (2008); the approach is unified and discussed in detail in Doberkat (2009b). The uniform assumption on an analytic base space makes many constructions possible for which this kind of topological scenario is essential, one prominent example being Souslin's Theorem, which gives, amongst other things, convenient conditions on a set's measurability – see Theorem 3.14. In short, since measure theory on analytic spaces is well developed, this is an attractive choice. On the other hand, one may ask whether these requirements are really necessary, that is, whether it is possible to investigate questions of expressivity without a topological scenario. Both of these issues are discussed in this paper.

We will begin by laying the foundations by discussing modal logics and Kripke models, first from the classical, relational point of view, then taking probabilities into account. We do this in Section 2, where we also introduce coalgebras and relate them to Krikpe models, and present a first discussion of bisimilarity. Since we will need some probability theory, we look at the probabilist's toolbox in Section 3, and lay out the definitions and tools we will need later. Here we also introduce the special brand of equivalence relations that we will need, and introduce the subprobability functor as one of the essential ingredients.

The step from modal logics to coalgebraic logics is taken in Section 4. We motivate the transition from modal operators to predicate transformations by introducing suitable natural transformations for the functors involved, amongst which we must always include the subprobability functor since we want the coalgebras to be stochastic in nature. This is in contrast to the usual approach, which is almost always based on functors on the category of sets. The emerging picture is far from being as clear cut as in the set based case; this becomes apparent when we contrast the two scenarios in Section 4.5. Predicate liftings are investigated further in Section 5 with the goal of determining the expressivity

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

in the interplay between the logics and their Kripke models. Although important tools from measure theory are missing, it is shown that a substantial number of the results obtained for analytic spaces can be carried over, sometimes by adapting techniques that were originally developed for the analytic case. An important technical property arises here, as in the set-based case, through the possibility of separating states through a formula – in this sense there have to be enough predicate liftings. Separation has an immediate consequence for the equivalence relation defined by the logic: it turns out to be a congruence, which is very useful for constructing new models from given ones. Our investigation of these properties is accompanied by examples dealing with some well-studied logics in Section 5.6.

Section 6 deals with bisimilarity, which is best expressed in terms of stochastic relations. A coalgebra for a functor $\mathfrak{S} \cdot \mathfrak{F}$ with $\mathfrak{F}$ an endofunctor on the category of measurable spaces, or one of its more prominent full subcategories such as Polish or analytic spaces, is a particular instance of a stochastic relation, so it seems sensible to discuss the problem in terms of stochastic relations. We introduce measurable selections and show that the semi-pullback for stochastic relations exists under certain conditions. Semi-pullbacks in turn are used to show that behaviourally equivalent relations are bisimilar. The corresponding result on selections is established by the selection theorem due to Kuratowski and Ryll-Nardzewski, we correct an argument used by the first author in Doberkat (2005). We indicate that selection arguments are a powerful tool for investigating bisimulations by relating congruences and bisimulations, and constructing a bridge to the usual coalgebraic interpretation of bisimulations for investigating coalgebras (Rutten 2000); the power of this line of argument was recently demonstrated in the identification of the quotient structure for the Giry monad through selection arguments (Doberkat 2010b). Finally, we relate the factor spaces of behaviourally equivalent relations by showing that they are isomorphic, and *vice versa*. These properties are related to bisimilarity. Note that the central arguments for investigating bisimilarity (measurable selections, Souslin's Theorem) are available only in a topological context; this makes bisimilarity a concept that is brought to life only in the presence of suitable topological assumptions.

Section 7 provides conclusions, mentions omitted topics and gives hints for further development.

2. Modal logics

In this section we give a very short introduction to modal logics by giving the general syntax of formulas, defining Kripke models and showing how these models interpret a modal logic (Section 2.1) – the exposition follows Blackburn *et al.* (2001). Kripke models are based on relations, but they can be reformulated from a coalgebraic perspective, which we do in Section 2.2. In this vein, morphisms are introduced as morphisms for the underlying coalgebra, indicating the more general perspective that we will focus on in this paper. We also introduce bisimulations through the corresponding relations, and in Section 2.3 we derive the Hennessy–Milner Theorem on the equivalence of bisimilarity and modal equivalence.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

2.1. *Definitions*

Let Var be a countable set of propositional letters, which is fixed throughout, and O be a set of modal operators. We assume that O is non-empty. $\tau = (O, \mathsf{ar})$ is said to be a *modal similarity type* if and only if $\mathsf{ar} : O \to \mathbb{N}$ is a map assigning each modal operator $\triangle$ its arity $\mathsf{ar}(\triangle) \geqslant 0$. The similarity type τ will be fixed.

We define the *modal language* $\mathscr{L}_{\neg}(\tau)$ based on τ and Var. The formulas are given by the syntax

$$\phi ::= \ p \ \mid \ \top \ \mid \ \phi_1 \wedge \phi_2 \ \mid \ \neg\phi \ \mid \ \triangle(\phi_1, \ldots \phi_{\mathsf{ar}(\triangle)}),$$

where $p \in \mathsf{Var}$ is a propositional letter. If we have $O = \{\Diamond\}$ with $\mathsf{ar}(\Diamond) = 1$, we obtain the formulas of the well-known *basic modal language* with negation.

This definition can be modified according to the application one has in mind: omitting negation in the basic modal language defines the formulas in the *negation-free modal language* $\mathscr{L}(\tau)$, the syntax of which should be evident.

The *extended modal language* is defined through the syntax

$$\phi ::= \ p \ \mid \ \top \ \mid \ \phi_1 \wedge \phi_2 \ \mid \ \neg\phi \ \mid \ \triangle_q(\phi_1, \ldots \phi_{\mathsf{ar}(\triangle)}),$$

where $q \in \mathbb{Q} \cap [0,1]$ is a rational number and $p \in \mathsf{Var}$ is a propositional letter. This can be modelled through the similarity type $\tau = \big(\{\,(\triangle, q) \mid \triangle \in O, q \in \mathbb{Q} \cap [0,1]\,\}, \mathsf{ar}'\big)$ with $\mathsf{ar}'(\triangle, q) := \mathsf{ar}(\triangle)$ for all q. From now on we will write $\triangle_q$ rather than $(\triangle, q)$. Again, if we have $O = \{\Diamond\}$ as the similarity type, we get an entirely new line of formulas through $(\Diamond_q)_{q \in \mathbb{Q} \cap [0,1]}$. This language is of interest when interpreting modal logics through stochastic Kripke models, as we will see below.

A *non-deterministic τ-Kripke model* $\mathsf{S} = (S, \gamma_\tau, V)$ consists of a state space S, a family $\gamma_\tau = (\gamma_\triangle)_{\triangle \in O}$ of set-valued maps $\gamma_\triangle : S \to \mathfrak{P}\left(S^{\mathsf{ar}(\triangle)}\right)$ and a set-valued map $V : \mathsf{Var} \to \mathfrak{P}(S)$. Here and in the following, we will use $\mathfrak{P}(S)$ to denote the powerset of a set S, thus $\mathfrak{P}(S) = \{A \mid A \subseteq S\}$.

The satisfaction relation $\models$ for a non-deterministic τ-Kripke model S is defined as usual for the basic modal language:

$$\begin{aligned}
\mathsf{S}, s \models p &\iff s \in V(p) \\
\mathsf{S}, s \models \neg\phi &\iff \mathsf{S}, s \not\models \phi \\
\mathsf{S}, s \models \phi_1 \wedge \phi_2 &\iff \mathsf{S}, s \models \phi_1 \text{ and } \mathsf{S}, s \models \phi_2 \\
\mathsf{S}, s \models \triangle(\phi_1, \ldots, \phi_{\mathsf{ar}(\triangle)}) &\iff \exists (s_1, \ldots, s_{\mathsf{ar}(\triangle)}) \in \gamma_\triangle(s) : \mathsf{S}, s_i \models \phi_i \\
&\qquad\qquad \text{for } 1 \leqslant i \leqslant \mathsf{ar}(\triangle).
\end{aligned}$$

We use

$$[\![\phi]\!]_{\mathsf{S}} := \{\, s \in S \mid \mathsf{S}, s \models \phi \,\}$$

to denote the set of states for which the formula ϕ is valid (the *extension of formula* ϕ), and

$$\mathrm{Th}_{\mathsf{S}}(s) := \{\, \phi \mid \mathsf{S}, s \models \phi \,\}$$

to denote the *theory of state* s in $\mathscr{S}$.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Note that functions $\gamma : S \to \mathfrak{P}(S^n)$ are in one-to-one correspondence with $(n+1)$-ary relations $R_\gamma \subseteq S^{n+1}$ through

$$(s_0, s_1, \ldots s_n) \in R_\gamma \iff (s_1, \ldots, s_n) \in \gamma(s_0).$$

Thus, we may consider the functions γ_Δ of a τ-Kripke model as successor relations. This point of view is popular in modal logic (Blackburn *et al.* 2001).

2.2. *Kripke models as coalgebras*

For simplicity, we assume that we have only one modal operator $\Diamond$ of arity one. Thus a Kripke model is interpreted as a triple (S, γ, V) with $\gamma : S \to \mathfrak{P}(S)$ and $V : \mathsf{Var} \to \mathfrak{P}(S)$. By dropping the valuations from the definition of Kripke models, we arrive at Kripke frames: a *Kripke frame* is given by a set S together with a function $\gamma : S \to \mathfrak{P}(S)$. Noting that $\mathfrak{P}$ in an endofunctor on the category **Set** of all sets and with functions as morphisms, a Kripke frame is interpreted as a coalgebra for the powerset functor $\mathfrak{P}$; *cf.* Section 3.5.

This interpretation is fairly natural, as we will see when studying morphisms. We define morphisms between Kripke frames as morphisms for the associated coalgebras. Morphisms of Kripke models are also supposed to preserve the valuations.

Definition 2.1. Let $\mathsf{S} = (S, \gamma, V)$ and $\mathsf{T} = (T, \delta, W)$ be Kripke models. A map $f : S \to T$ is said to be a *morphism* $\mathsf{S} \to \mathsf{T}$ if and only if $f^{-1}[W] = V$, and the following diagram commutes:

$$\begin{array}{ccc} S & \xrightarrow{\quad f \quad} & T \\ {\scriptstyle\gamma}\downarrow & & \downarrow{\scriptstyle\delta} \\ \mathfrak{P}(S) & \xrightarrow{\quad \mathfrak{P}(f) \quad} & \mathfrak{P}(T) \end{array}$$

Here we have written $\mathfrak{P}(f) : \mathfrak{P}(S) \to \mathfrak{P}(T)$ for the image-function of the function $f : S \to T$, so

$$\mathfrak{P}(f)(A) := f[A] = \{ f(a) \mid a \in A \}$$

for $A \in \mathfrak{P}(S)$. Thus we have for a morphism $\mathsf{S} \to \mathsf{T}$ for each state $s \in S$

$$\delta(f(s)) = \mathfrak{P}(f)(\gamma(s)) = f[\gamma(s)],$$

so $t \in \delta(f(s))$ if and only if there exists $s' \in \gamma(s)$ such that $t = f(s')$.

Lemma 2.2. Let $\mathsf{S} = (S, \gamma, V)$ and $\mathsf{T} = (T, \delta, W)$ be Kripke models. A map $f : S \to T$ is a morphism $\mathsf{S} \to \mathsf{T}$ if and only if:

(a) $f^{-1}[W(p)] = V(p)$ for all $p \in \mathsf{Var}$.
(b) $f(s') \in \delta(f(s))$ for all $s' \in \gamma(s)$.
(c) If $t \in \delta(f(s))$, there exists $s' \in \gamma(s)$ such that $t = f(s')$.

Condition (b) is usually called the *forward condition*: given a state in $\gamma(s)$, its image is a member of the transitions associated with $f(s)$. Correspondingly, condition (c) is called the *backward condition*: each transition associated with $f(s)$ may be obtained as the image

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

of a transition associated with state s under f. But this is just the classical definition of a bounded morphism (Blackburn *et al.* 2001, Definition 2.12) in modal logics.

An immediate consequence of Lemma 2.2 is that morphisms respect and reflect the validity of formulas.

Corollary 2.3. Let $\mathsf{S} = (S, \gamma, V)$ and $\mathsf{T} = (T, \delta, W)$ be Kripke models and $f : \mathsf{S} \to \mathsf{T}$ be a morphism. Then

$$\mathsf{S}, s \models \phi \iff \mathsf{T}, f(s) \models \phi$$

holds for each state $s \in S$ and for each formula ϕ. That is, $\mathrm{Th}_{\mathsf{S}}(s) = \mathrm{Th}_{\mathsf{T}}(f(s))$ holds for each state $s \in S$.

Proof. The proof is by induction on the structure of formula ϕ. □

2.3. *Bisimilar Kripke models*

Bisimilarity is a fundamental property used in the study of modal logics and their Kripke models.

Definition 2.4. Let $\mathsf{S} = (S, \gamma, V)$ and $\mathsf{T} = (T, \delta, W)$ be Kripke models.

(1) A relation $Z \subseteq S \times T$ is said to be a *bisimulation between* S *and* T if and only if

(i) If $(s, t) \in Z$, s and t satisfy the same propositional letters.

(ii) If $(s, t) \in Z$ and $s' \in \gamma(s)$, there exists $t' \in \delta(t)$ with $(s', t') \in Z$ (the *forward* condition).

(iii) If $(s, t) \in Z$ and $t' \in \delta(t)$, there exists $s' \in \gamma(s)$ with $(s', t') \in Z$ (the *backward* condition).

(2) States $s \in S$ and $t \in T$ are said to be *bisimilar* if and only if there exists a bisimulation between S and T containing (s, t).

Bisimulations between Kripke frames (S, γ) and (T, δ) are defined to be those relations $Z \subseteq S \times T$ that satisfy (ii) and (iii) of Part (1).

Thus states are bisimilar if and only if they have identical information with respect to the propositional letters, and if the transition behaviour is reflected in the relation: roughly speaking, bisimilar states move into bisimilar states.

The following lemma is an easy exercise (Rutten 2000) relating bisimulations and morphisms of Kripke models.

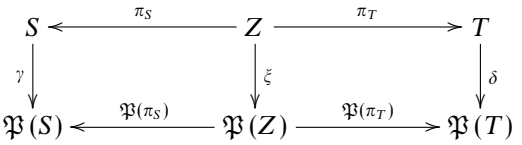
Lemma 2.5. Let $\mathsf{S} = (S, \gamma, V)$ and $\mathsf{T} = (T, \delta, W)$ be Kripke models. A relation $Z \subseteq S \times T$ is a bisimulation between S and T if and only if we can find $\xi : Z \to \mathfrak{P}(Z)$ and $U : \mathsf{Var} \to \mathfrak{P}(Z)$ such that the evident projections $\pi_S : Z \to S$ and $\pi_T : Z \to T$ become morphisms of Kripke models $(Z, \xi, U) \to \mathsf{S}$ and $(Z, \xi, U) \to \mathsf{T}$, respectively.

Spelling out the conditions in Lemma 2.5, we see that we are forced to set

$$U(p) := (V(p) \times W(p)) \cap Z,$$

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

and that we require that the map $\xi : Z \to \mathfrak{P}(Z)$ makes the following diagram commute:



Modal formulas are invariant under bisimulation: bisimilar states have the same theories.

Proposition 2.6. Let $\mathsf{S} = (S, \gamma, V)$ and $\mathsf{T} = (T, \delta, W)$ be Kripke models, and assume that $s \in S$ and $t \in T$ are bisimilar. Then $\mathrm{Th}_{\mathsf{S}}(s) = \mathrm{Th}_{\mathsf{T}}(t)$.

Proof. Let Z be a bisimulation between S and T with $(s,t) \in Z$. By Lemma 2.5, we find a Kripke model $\mathsf{Z} = (Z, \xi, U)$ such that the projections $\pi_S : \mathsf{Z} \to \mathsf{S}$ and $\pi_T : \mathsf{Z} \to \mathsf{T}$ are morphisms. Thus, the statement follows from Corollary 2.3. □

A Kripke model (S, γ, V) is said to be *image finite* if and only if $\gamma(s)$ is finite for each $s \in S$. For these image finite Kripke models, *modal equivalence*, that is, the equality of theories, yields a bisimulation.

Lemma 2.7. Let $\mathsf{S} = (S, \gamma, V)$ and $\mathsf{T} = (T, \delta, W)$ be image finite Kripke models. Then

$$\{\, (s,t) \in S \times T \mid \mathrm{Th}_{\mathsf{S}}(s) = \mathrm{Th}_{\mathsf{T}}(t) \,\}$$

is a bisimulation.

Proof. We will just concentrate on the forward condition in Definition 2.4; the backward condition is established in the same way, *mutatis mutandis*. Assume $\mathrm{Th}_{\mathsf{S}}(s) = \mathrm{Th}_{\mathsf{T}}(t)$ and $s' \in \gamma(s)$. So we have to find $t' \in \delta(t)$ such that s' and t' have the same theories. Assuming the contrary, we would have $\delta(t) \neq \varnothing$, because otherwise $\mathsf{T}, t \models \neg\Diamond\top$, but $\mathsf{S}, s \models \Diamond\top$ since $\gamma(s) \neq \varnothing$, which contradicts $\mathrm{Th}_{\mathsf{S}}(s) = \mathrm{Th}_{\mathsf{T}}(t)$. Thus $\delta(t) = \{t_1, \ldots, t_n\}$ for some $n > 0$. We can find for each j a formula ϕ_j such that $\mathsf{S}, s' \models \phi_j$ but $\mathscr{T}, t_j \not\models \phi_j$. Consequently, $\mathsf{S}, s \models \Diamond(\phi_1 \wedge \cdots \wedge \phi_n)$ and $\mathsf{T}, t \not\models \Diamond(\phi_1 \wedge \cdots \wedge \phi_n)$, which contradicts the assumption $\mathrm{Th}_{\mathsf{S}}(s) = \mathrm{Th}_{\mathsf{T}}(t)$ again. □

This observation yields the Hennessy–Milner Theorem as an interesting consequence, *viz.* that bisimilarity and modal equivalence are the same for image finite models.

Theorem 2.8. (Hennessy–Milner) Let $\mathsf{S} = (S, \gamma, V)$ and $\mathsf{T} = (T, \delta, W)$ be image finite Kripke models, then states $s \in S$ and $t \in T$ are modal equivalent if and only if they are bisimilar.

3. Measure theoretic and other preliminaries

In this section we recall the necessary concepts and results from measure theory: standard references include Halmos (1950), Rudin (1974), Dunford and Schwartz (1957) and Billingsley (1995); special emphasis on Polish spaces is given in Parthasarathy (1967) and

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Billingsley (1968); and for a tutorial treatment geared towards Markov transition systems and modal logics, see Doberkat (2007b).

3.1. *Measurable spaces*

A measurable space X consists of a set $|X|$ and a σ-algebra $\mathfrak{B}X$ on X, that is, a family of subsets of $|X|$ that is closed under complementation, countable intersections and countable unions. For each family $\mathscr{A}$ of subsets of a set M there is a smallest σ-algebra on M containing $\mathscr{A}$, which we denote by $\sigma(\mathscr{A})$. If $\mathfrak{B}X = \sigma(\mathscr{A})$, then $\mathscr{A}$ is said to be a *generator* of $\mathfrak{B}X$. A *π-class* is a family of subsets that is closed under finite intersections. In particular, every π-class on a set M contains M. A generator that is also a π-class is said to be a *π-generator*. A family $\mathscr{A}$ of subsets of a set M is said to be a *λ-class* if $\mathscr{A}$ is closed under complementation and countable disjoint unions. Of course, every σ-algebra is a λ-class.

The importance of π and λ-classes for our purposes stems from the following well-known π-λ-theorem (Kechris 1994, Theorem 10.1).

Theorem 3.1. Let $\mathscr{A}$ be a π-class of subsets of a set M. Then $\sigma(\mathscr{A})$ is the smallest λ-class containing $\mathscr{A}$.

A *measurable function* $X \to Y$ is given by a function $f : |X| \to |Y|$ such that $f^{-1}[B] \in \mathfrak{B}X$ for all $B \in \mathfrak{B}Y$. When $\mathfrak{B}Y = \sigma(\mathscr{A})$, measurability of f is guaranteed by $f^{-1}[A] \in \mathfrak{B}X$ for all $A \in \mathscr{A}$. The category of measurable spaces with measurable functions as morphisms is denoted by **Meas**. Observe that we do not distinguish notationally between a **Meas**-morphism and its underlying function, and we will often just write X in place of $|X|$.

The assignment $X \mapsto \mathfrak{B}X$ defines a functor $\mathbf{Meas}^{\mathrm{op}} \to \mathbf{Set}$ with its action on morphisms given by the restriction of the inverse-image function, thus $\mathfrak{B}f = f^{-1} : \mathfrak{B}Y \to \mathfrak{B}X$ for $f : X \to Y$. Observe that we have a natural isomorphism

$$\mathfrak{B} \cong \mathbf{Meas}(-, \mathbf{2}), \tag{2}$$

with $\mathbf{2}$ the two-point space in which every subset is measurable. In particular, we obtain the following lemma.

Lemma 3.2. $\mathfrak{B} : \mathbf{Meas}^{\mathrm{op}} \to \mathbf{Set}$ preserves limits and $\mathfrak{B}f$ is injective provided f is surjective.

Proof. The first statement follows from the fact that every hom-functor preserves limits (Mac Lane 1997, Theorem V.4.1). The proof of the second statement is obvious. □

3.2. *Measurable morphisms*

We will now turn our attention to properties of the category **Meas** of measurable spaces.

Given a family $(Y_i)_{i\in I}$ of measurable spaces and a family $(f_i : A \to |Y_i|)_{i\in I}$ of functions with common domain, we define the *initial σ-algebra* with respect to $(f_i)_{i\in I}$ by

$$\mathscr{A} = \sigma\Big(\bigcup_I f_i^{-1}[\mathfrak{B}Y_i]\Big). \tag{3}$$

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

This σ-algebra is characterised by the following property: a function $g : |X| \to A$ is measurable with respect to $\mathfrak{B}X$ and $\mathscr{A}$ if and only if all $f_i \cdot g : X \to Y_i$ are measurable. For a one-element family $I = \{*\}$, the expression (3) simplifies to $\mathscr{A} = f_*^{-1}[\mathfrak{B}(Y_*)]$. Accordingly, we say a morphism $f : X \to Y$ is *initial* if $\mathfrak{B}X = f^{-1}[\mathfrak{B}Y]$.

Lemma 3.3. Let $(f_i : A \to |Y_i|)_{i\in I}$ be a family of functions and assume that each Y_i has a generator $\mathscr{G}_i$. Then the initial σ-algebra with respect to $(f_i)_{i\in I}$ is $\sigma\left(\bigcup_I f_i^{-1}[\mathscr{G}_i]\right)$.

Proof. See Elstrodt (1999, Satz 5.2). □

Dually, the *final σ-algebra* with respect to $(g_i : |X_i| \to B)_{i\in I}$ is $\{\, E \subseteq B \mid \forall i \in I : g_i^{-1}[E] \in \mathfrak{B}X_i \,\}$. It is characterised by the property that measurability of $h : B \to |Y|$ is guaranteed by the measurability of all $h \cdot g_i$. We say a morphism $f : X \to Y$ is *final* if $\mathfrak{B}Y$ is the final σ-algebra with respect to f, $\mathfrak{B}X$.

This can be used to show that **Meas** is complete and cocomplete. Limits are constructed as follows: construct the limit $\left(L, (\ell_i)\right)_{i\in I}$ of the underlying diagram in **Set** and equip L with the initial σ-algebra with respect to the projections ℓ_i, and thus with the σ-algebra generated by $\bigcup \ell_i^{-1}[\mathfrak{B}X_i]$. Colimits are formed dually using final σ-algebras.

Thus, for instance, the coproduct X of a family $(X_i)_{i\in I}$ of measurable spaces is constructed by taking as its underlying set the disjoint union of the underlying sets of the X_i and declaring $A \subseteq |X|$ to be measurable if and only if $A \cap |X_i| \in \mathfrak{B}X_i$ holds for all $i \in I$.

3.2.1. *Point-separating measurable spaces.* A family $\mathscr{A}$ of subsets of a set X is said to *separate points* if whenever $x \neq x'$ there exists $A \in \mathscr{A}$ with $x \in A$, $x' \notin A$, or *vice versa.*

Lemma 3.4. $\mathscr{A}$ separates points if and only if $\sigma(\mathscr{A})$ separates points.

Definition 3.5. A measurable space X is said to be *point-separating* if $\mathfrak{B}X$ separates points.

A family $(f_i : X \to Y_i)_{i\in I}$ of morphisms with common domain in an arbitrary category **C** is said to be *jointly monic* if, for every pair $g, h : W \to X$ of morphisms, we have that g and h are equal provided $f_i \cdot g = f_i \cdot h$ holds for all $i \in I$. In **Set** or **Meas**, a family $(f_i)_{i\in I}$ as above is jointly monic if and only if $x \neq x'$ in X implies that there exists some $i \in I$ with $f_i(x) \neq f_i(x')$. It is well known that the projections $(\ell_i)_{i\in I}$ of a limit are always jointly monic.

Proposition 3.6. If $(f_i : X \to Y_i)_{i\in I}$ is a jointly monic family of morphisms in **Meas** and each Y_i is point-separating, then X is point-separating also. Moreover, the coproduct of a family of point-separating measurable spaces is point-separating.

Proof. We will just consider the first claim. Let $(f_i : X \to Y_i)_{i\in I}$ be jointly monic in **Meas** and assume that each Y_i is point-separating. Thus, for $x \neq x'$ in X, we find an $i \in I$ with $f_i(x) \neq f_i(x')$ in Y_i, and hence, by assumption, $B \in \mathfrak{B}Y_i$ with, say, $f_i(x) \in B$ and $f_i(x') \notin B$. Since f_i is measurable, we have $A := f_i^{-1}[B] \in \mathfrak{B}X$ and $x \in A$, $x' \notin A$ holds. □

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

3.2.2. *Equivalence relations and invariant sets.* For an equivalence relation α on a set X, we use $\eta_\alpha : X \to X/\alpha$ to denote the projection sending an element x to its equivalence class $[x]_\alpha$. Observe that $\eta_\alpha\left[\eta_\alpha^{-1}[B]\right] = B$ holds for every $B \subseteq X/\alpha$ because η_α is surjective. The following result is elementary.

Lemma 3.7. The following statements are equivalent for $A \subseteq X$:

(1) A is α-*invariant*, that is, $x \in A$ and $x \,\alpha\, x'$ imply $x' \in A$ for all $x, x' \in X$.
(2) A is a union of equivalence classes. Thus, $A = \bigcup\{\, [x]_\alpha \mid x \in A \,\}$ holds.
(3) $\eta_\alpha^{-1}\,[\eta_\alpha[A]] = A$ holds.

Now suppose that X is actually a measurable space. We write X/α for the measurable space whose underlying set is $|X|/\alpha$ and whose σ-algebra is the final σ-algebra with respect to $\mathfrak{B}X$ and η_α. Thus, a subset $B \subseteq X/\alpha$ is measurable if and only if $\eta_\alpha^{-1}[B] \in \mathfrak{B}X$ holds. Thus, by definition, we have a pair of mutually inverse bijections

$$\Sigma(\mathfrak{B}X, \alpha) \underset{\eta_\alpha^{-1}[-]}{\overset{\eta_\alpha[-]}{\rightleftarrows}} \mathfrak{B}(X/\alpha),$$

where $\Sigma(\mathfrak{B}X, \alpha) = \{\, A \in \mathfrak{B}X \mid A \text{ is } \alpha\text{-invariant} \,\}$. That is, the α-invariant subsets provide a useful description of the measurable subsets of the factor space in terms of subsets of the original space. It is not difficult to see that $\Sigma(\mathfrak{B}X, \alpha)$ is in fact a σ-algebra.

Each set $\mathscr{A}$ of subsets on a set M induces an equivalence relation $\approx_{\mathscr{A}}$ on M upon setting

$$x \approx_{\mathscr{A}} x' \text{ if and only if } \forall A \in \mathscr{A} : [x \in A \iff x' \in A].$$

Then each $A \in \mathscr{A}$ is $\approx_{\mathscr{A}}$-invariant. It is then not difficult to see that

$$\approx_{\mathscr{A}} = \approx_{\sigma(\mathscr{A})},$$

holds, *cf.* Srivastava (1998, Lemma 3.1.6). We will need this result later.

3.3. *Subprobability functor*

Our main example of a functor on measurable spaces will be the subprobability functor, which we are going to introduce now.

Fix a measurable space X. A *subprobability on* X is a function $\mu : \mathfrak{B}X \to [0, 1]$ satisfying

$$\mu\Big(\bigcup_{i \in \mathbb{N}} A_i\Big) = \sum_{i \in \mathbb{N}} \mu(A_i)$$

for every countable family (A_i) of pairwise disjoint sets in $\mathfrak{B}X$. Note that this implies $\mu(\varnothing) = 0$, and that μ is monotone.

The set of all subprobability measures on X becomes a measurable space $\mathfrak{S}X$ when equipped with the *weak-*-σ-algebra*, which is defined as the initial σ-algebra with respect to $(\mathsf{ev}_A)_{A \in \mathfrak{B}X}$ with $\mathsf{ev}_A : \mathfrak{S}X \to [0, 1]$, $\mu \mapsto \mu(A)$. Here, $[0, 1]$ is equipped with the Borel sets, which is defined as the σ-algebra that is generated by the sets of the form $[q, 1]$, with $q \in \mathbb{Q} \cap [0, 1]$.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

This subprobability construction gives rise to a functor $\mathfrak{S} : \mathbf{Meas} \to \mathbf{Meas}$ by setting

$$\mathfrak{S}f(\mu)(B) = \mu(f^{-1}[B])$$

for $f : X \to Y$ in **Meas**, $\mu \in \mathfrak{S}X$, $B \in \mathfrak{B}Y$.

Lemma 3.8. Let X be a measurable space and $\mathscr{G} \subseteq \mathfrak{B}X$ be a π-generator. Then $\mathfrak{B}(\mathfrak{S}X)$ is generated by the set

$$\{\, \mathsf{ev}_G^{-1}[q,1] \mid G \in \mathscr{G}, q \in \mathbb{Q} \cap [0,1] \,\}.$$

Moreover, $\mu, \nu \in \mathfrak{S}X$ are equal if they agree on $\mathscr{G}$.

The proof we will present is taken from Jacobs and Sokolova (2010); the result was also proved, but using different methods, in Viglizzo (2005a).

Proof.

(1) We write

$$\mathscr{D} = \{\, \mathsf{ev}_G^{-1}[q,1] \mid G \in \mathscr{G}, q \in \mathbb{Q} \cap [0,1] \,\}.$$

It is clear that $\mathscr{D} \subseteq \mathfrak{B}(\mathfrak{S}X)$, hence $\sigma(\mathscr{D}) \subseteq \mathfrak{B}(\mathfrak{S}X)$. Our goal is to show the opposite inclusion. We write

$$\mathscr{Q} := \{\, [q,1] \mid q \in \mathbb{Q} \,\} \subseteq \mathfrak{B}([0,1])$$

and consider the set

$$\mathscr{C} := \{\, C \in \mathfrak{B}X \mid \forall Q \in \mathscr{Q} : \mathsf{ev}_C^{-1}[Q] \in \sigma(\mathscr{D}) \,\}.$$

Since $\sigma(\mathscr{Q})$ is the usual Borel structure on $[0,1]$, we may employ results from the theory of measurable functions $X \to \mathbb{R}$. Observe that $\mathscr{G} \subseteq \mathscr{C}$ holds by construction.

(2) We will show that $\mathscr{C}$ is a λ-class. Take any $C \in \mathscr{C}$. We have $\mathsf{ev}_{X \setminus C} = \mathsf{ev}_X - \mathsf{ev}_C$ pointwise. Since the set of measurable functions is a vector-space under pointwise addition and scalar-multiplication, we obtain $X \setminus C \in \mathscr{C}$ if $C \in \mathscr{C}$; thus, $\mathscr{C}$ is closed under complementation.

We now take a sequence $(C_i)_{i \in \mathbb{N}}$ of pairwise disjoint elements of $\mathscr{C}$ and write $C = \bigcup C_i$. σ-additivity of subprobability measures given $\mathsf{ev}_C = \sum_i \mathsf{ev}_{C_i}$. The sequence (f_k) with $f_k = \sum_{i=0}^{k} \mathsf{ev}_{C_i}$ converges pointwise to ev_C and is increasing. Since each f_k is measurable, so is the limit ev_C. Hence $\mathscr{C}$ is closed under countable disjoint unions.

Since $\mathscr{G}$ is a π-generator of $\mathfrak{B}X$, $\mathfrak{B}X$ is the smallest λ-class containing $\mathscr{G}$ by Theorem 3.1. Thus $\mathscr{C} = \mathfrak{B}X$, and hence $\mathfrak{B}(\mathfrak{S}X) \subseteq \sigma(\mathscr{D})$.

(3) The second claim is shown using the same methods. Consider the set

$$\mathscr{C} = \{\, C \in \mathfrak{B}X \mid \mu(C) = \nu(C) \,\}.$$

This contains $\mathscr{G}$ and is easily shown to be a λ-class. □

Lemma 3.9. Each $\mathfrak{S}X$ is point-separating.

Proof. Since, given $\mu, \mu' \in \mathfrak{S}X$ with $\mu \neq \mu'$, there exists $A \in \mathfrak{B}X$ with $\mathsf{ev}_A(\mu) \neq \mathsf{ev}_A(\mu')$, we conclude that $(\mathsf{ev}_A)_{A \in \mathfrak{B}X}$ is jointly monic and the claim follows from Proposition 3.6. □

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

3.3.1. *Stochastic relations and stochastic Kripke models.* We will now define a stochastic counterpart to the set-valued relations that non-deterministic Kripke models are based on.

Definition 3.10. A *stochastic relation* $K : X \rightsquigarrow Y$ between measurable spaces X and Y is a measurable function $X \to \mathfrak{S}Y$.

Unwinding the condition of measurability on a stochastic relation, we may say that a function $K : X \to |\mathfrak{S}Y|$ is a stochastic relation if and only if the function $K(-)(B) : X \to [0,1]$ is measurable for each $B \in \mathfrak{B}Y$.

Now that stochastic relations are available, we can define a probabilistic counterpart of non-deterministic Kripke models. Let $\tau = (O, \mathsf{ar})$ be a modal similarity type, see Section 2.1, and Var a countable set of variables. A *stochastic τ-Kripke model* $\mathsf{K} = (X, K_\tau, V)$ has a state space X that is a measurables space, a family $K_\tau = (K_\triangle)_{\triangle \in O}$ of stochastic relations $K_\triangle : X \rightsquigarrow X^{\mathsf{ar}(\triangle)}$, and a set-valued map $V : \mathsf{Var} \to \mathfrak{B}X$.

The interpretation of formulas of the extended modal language in a stochastic τ-Kripke model K follows the interpretation for non-deterministic Kripke models given in Section 2.1, the only difference arises when a modal operator is involved:

$$\mathsf{K}, x \models \triangle_q(\phi_1, \ldots, \phi_{\mathsf{ar}(\triangle)})$$

holds if and only if there exist measurable subsets $A_1, \ldots, A_{\mathsf{ar}(\triangle)}$ of X such that $\mathsf{K}, x_i \models \phi_i$ holds for all $x_i \in A_i$ for $1 \leqslant i \leqslant \mathsf{ar}(\triangle)$, and

$$K_\triangle(x)(A_1 \times \ldots \times A_{\mathsf{ar}(\triangle)}) \geqslant q.$$

With state transition systems in mind, this interpretation of validity reflects the fact that upon the move indicated by $\triangle_q$, a state x satisfies $\triangle_q(\phi_1, \ldots, \phi_{\mathsf{ar}(\triangle)})$ if and only if we can find states x_i satisfying ϕ_i with a $K_\triangle$-probability not smaller than q. Note that the usual operators $\triangle$ and their dual modal operators ∇ are replaced by a whole spectrum of operators $\triangle_q$ that permit a finer and probabilistically more adequate notion of satisfaction.

Again, let $[\![\phi]\!]_\mathsf{K}$ be the set of all states for which formula ϕ is satisfied under K, and

$$\mathrm{Th}_\mathsf{K}(x) := \{\, \phi \mid \mathsf{K}, x \models \phi \,\}$$

the theory for state $x \in X$.

An easy inductive argument shows that the sets $[\![\phi]\!]_\mathsf{K}$ are measurable, so they may be used as arguments for the stochastic relations we are working with.

Lemma 3.11. $[\![\phi]\!]_\mathsf{K}$ is a measurable subset of X for each ϕ in the extended modal language.

3.4. *A small dose of topology*

We need some topological preliminaries for the exploration of bisimilarity. Given a topology $\mathscr{T}$ on a set X, the smallest σ-algebra containing $\mathscr{T}$ is called the σ-algebra of *Borel sets*; it is usually denoted by $\mathscr{B}(X, \mathscr{T})$, or just by $\mathscr{B}(X)$ if the topology is understood. We usually omit the notation of the Borel sets for a topological space, and in that case

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

will always mean this σ-algebra. Similarly, Borel measurability here always refers to the Borel sets. A topological space $(X, \mathscr{T})$ is said to be *Polish* if and only if there exists a metric d for topology $\mathscr{T}$ such that (X, d) is a complete metric space with a countable dense subset. Polish spaces have some pleasant measure theoretic properties, as we will see. An *analytic space* is the image of a continuous function between Polish spaces. A *Standard Borel space* is a measurable space $(X, \mathscr{A})$ such that we can find a Polish topology $\mathscr{T}$ on X with $\mathscr{B}(X, \mathscr{T}) = \mathscr{A}$.

First note that we can manipulate a Polish topology in such a way that Borel measurable maps into separable metric spaces become continuous.

Proposition 3.12. Let $(Y, \mathscr{S})$ be a separable metric space with topology $\mathscr{S}$. If $f : X \to Y$ is a $\mathscr{B}(X, \mathscr{T})$-$\mathscr{B}(Y, \mathscr{S})$-Borel measurable map, then there exists a Polish topology $\mathscr{T}'$ on X with these properties:

(a) $\mathscr{T}'$ is finer than $\mathscr{T}$.
(b) $\mathscr{T}$ and $\mathscr{T}'$ have the same Borel sets.
(c) f is $\mathscr{T}'$-$\mathscr{S}$ continuous.

Proof. See Srivastava (1998, Corollary 3.2.6) or Doberkat (2007b, Proposition 1.28). □

This property is technically most useful as it allows us to make measurable maps continuous when they go into a separable or second countable metric space (thus, in particular, into a Polish space). Consequently, we can capitalise on the properties of continuous maps for Borel maps.

The topological structure of Polish spaces is completely known through the famous characterisation of Polish spaces due to Alexandrov (Kuratowski 1966, III.33.VI).

Theorem 3.13 (Alexandrov). Let X be a separable metric space. Then X is homeomorphic to a subspace of the Hilbert cube $[0, 1]^\infty$. If X is Polish, this subspace is the intersection of a countable number of open sets.

Alexandrov's Theorem presents a complete characterisation of Polish spaces as the G_δ-sets of the Hilbert cube. Souslin's Theorem is a characterisation of Borel sets in an analytic space (Kechris 1994, Exercise 28.3).

Theorem 3.14 (Souslin). Let X be an analytic space. Then $A \in \mathscr{B}(X)$ if and only if both A and $X \setminus A$ are analytic.

Souslin's Theorem will be very helpful when we come to characterise images of Borel sets under certain Borel maps.

The set $\mathfrak{S}(X) = \mathfrak{S}(X, \mathscr{B}(X))$ of all subprobability measures is equipped with the *topology of weak convergence* whenever (X, d) is a metric space. The topology of weak convergence is the smallest topology on $\mathfrak{S}(X)$ that makes the map $\mu \mapsto \int_X f \, d\mu$ continuous for each continuous and bounded $f : X \to \mathbb{R}$. Convergence with respect to the weak topology is denoted by $\rightharpoonup_w$.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

This topology may be metrised for a Polish space. Before we go on to define a particular metric, we will make sure that we are not going to leave the realm of Polish spaces (Parthasarathy 1967).

Theorem 3.15. Let X be a Polish space. Then $\mathfrak{S}(X)$ is a Polish space in the topology of weak convergence. If X is compact metric, so is $\mathfrak{S}(X)$ with this topology. □

The Hutchinson metric is one of the metrics that have been studied in the literature for metrising the weak topology. Before defining it, recall that for a metric space (X, d) a map $f : X \to \mathbb{R}$ is said to be *1-Lipschitz* if and only if

$$\forall x, x' \in X : |f(x) - f(x')| \leqslant d(x, x').$$

Clearly, a 1-Lipschitz map is uniformly continuous. We use $\mathbf{L}_1(X)$ to denote the set of all 1-Lipschitz maps f on X with $\sup_{x \in X} |f(x)| \leqslant 1$.

We define for $\mu, \mu' \in \mathfrak{S}(X)$ their *Hutchinson distance*

$$\mathfrak{d}_{\mathfrak{H}}(\mu, \mu') := \sup\Big\{ \Big| \int_X f \; d\mu - \int_X f \; d\mu' \Big| \mid f \in \mathbf{L}_1(X) \Big\}.$$

Hence, this metric is suitable for our purposes – see Section 6.1, in particular, Proposition 6.2.

Proposition 3.16. For a Polish space X, $\mathfrak{d}_{\mathfrak{H}}$ is a metric on $\mathfrak{S}(X)$ for the topology of weak convergence.

Proof. See Edgar (1998, Theorem 2.5.17). □

The Borel sets for this topology are well known (Kechris 1994, Theorem 17.24).

Proposition 3.17. Let X be a separable metric space. The collection of Borel sets for the topology of weak convergence on $\mathfrak{S}(X)$ is the weak*-σ-algebra introduced in Section 3.3.

The topology of weak convergence will be used in Section 6.1.

3.4.1. *Komolgorov consistency.* We will now show that the subprobability functor $\mathfrak{S}$ preserves certain limits of Standard Borel spaces. In general, the relationships between $\mathfrak{S}$ and limits is quite complex, *cf.* Section 6.3.

We will first recall a classical result from Measure Theory, the famous Kolmogorov Consistency Theorem; see, for example, Parthasarathy (1967, Chapter V.5).

Theorem 3.18 (Kolmogorov consistency). Let

$$A_0 \xleftarrow{f_0} A_1 \xleftarrow{f_1} A_2 \ldots A_n \xleftarrow{f_n} A_{n+1} \ldots$$

be a sequence of Standard Borel spaces and surjective measurable functions. If, for each $i \in \mathbb{N}$, we have μ_i as a subprobability measure on A_i such that $\mathfrak{S}f_i(\mu_{i+1}) = \mu_i$ holds, then there is a unique subprobability measure μ on the limit L of the above sequence such that $\mu_i = \mathfrak{S}p_i(\mu)$, with p_i the projection.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

We write **SB** for the category of Standard Borel spaces with measurable maps as morphisms. Note that **SB** is closed under countable limits and countable coproducts in **Meas**.

Corollary 3.19. The subprobability functor $\mathfrak{S}$ preserves limits of ω^{op}-chains of surjective **SB**-morphisms.

Proof.

(1) Let $(f_i : A_{i+1} \to A_i)_{i\in\omega}$ be a sequence of surjective **SB**-morphisms with limit $(L, (\ell_i))$. We claim that $(\mathfrak{S}L, (\mathfrak{S}\ell_i))$ is a limit of the sequence $(\mathfrak{S}f_i : \mathfrak{S}A_{i+1} \to \mathfrak{S}A_i)$. Let X be a measurable space and $(t_i : X \to \mathfrak{S}A_i)$ be a sequence of measurable functions with $\mathfrak{S}f_i \cdot t_{i+1} = t_i$ for all $i \in \omega$. By the Kolmogorov Consistency Theorem (Theorem 3.18), there is, for each $x \in X$, a unique measure μ_x on $\mathfrak{S}L$ with $\mathfrak{S}\ell_i(\mu_x) = t_i(x)$ for all $i \in \omega$. Hence, there is a unique function $k : X \to \mathfrak{S}L$ with $\mathfrak{S}\ell_i \cdot k = t_i$ for all $i \in \omega$.

(2) We now need to show that k is measurable. To do this, we show that k is measurable, we will first show that $\mathscr{A} = \bigcup_{i\in\omega} \ell_i^{-1}[\mathfrak{B}A_i]$ is a Boolean algebra of subsets of L. Clearly, L and $\varnothing$ are in $\mathscr{A}$. To show that $\mathscr{A}$ is closed under complements, take any $A = \ell_i^{-1}[B]$ for some $i \in \omega$, $B \in \mathfrak{B}(A_i)$. We have

$$L \setminus \ell_i^{-1}[B] = \ell_i^{-1}[A_i \setminus B],$$

and thus $L \setminus A \in \mathscr{A}$ since $\mathfrak{B}(A_i)$ is closed under complements. To show closedness under finite intersections, take $i, j \in \omega$ and $C \in \mathfrak{B}(A_i)$, $D \in \mathfrak{B}(A_j)$. We may assume $i \leqslant j$. We write f for the composite $f_{j-1} \cdot \cdots \cdot f_i : A_j \to A_i$, and when $i = j$, set $f = 1_{A_i}$. Since f is measurable and $f \cdot \ell_j = \ell_i$, we have

$$\begin{aligned} \ell_i^{-1}[C] \cap \ell_j^{-1}[D] &= (f \cdot \ell_j)^{-1}[C] \cap \ell_j^{-1}[D] \\ &= \ell_j^{-1} \cdot f^{-1}[C] \cap \ell_j^{-1}[D] \\ &= \ell_j^{-1}[f^{-1}[C] \cap D]. \end{aligned}$$

Since $E = f^{-1}[C] \cap D$ is in $\mathfrak{B}(A_j)$ by measurability of f, we obtain

$$\ell_i^{-1}[C] \cap \ell_j^{-1}[D] \in \mathscr{A}.$$

(3) Thus, by Lemma 3.8, it suffices to show that $\mathsf{ev}_B \cdot k : X \to [0,1]$ is measurable for each $B \in \bigcup_{i\in\omega} \ell_i^{-1}[\mathfrak{B}(A_i)]$. We fix $i \in \omega$ and $C \in \mathfrak{B}(A_i)$ and write $B = \ell_i^{-1}[C]$. We then observe that

$$\begin{aligned} \mathsf{ev}_B \cdot k(x) &= k(x)(B) \\ &= k(x)(\ell_i^{-1}[C]) \\ &= \mathfrak{S}\ell_i \cdot k(x)(C) \\ &= t_i(x)(C) \\ &= \mathsf{ev}_C \cdot t_i(x) \end{aligned}$$

for each $x \in X$. Hence, we obtain $\mathsf{ev}_B \cdot k = \mathsf{ev}_C \cdot t_i$, and the latter function is measurable. □

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

The restriction to Standard Borel spaces as well as to surjective functions is necessary: arbitrary limits of ω^{op}-chains in **Meas** are not preserved by $\mathfrak{S}$ – see Viglizzo (2005b) and Yamasaki (1985) for a discussion. It turns out that $\mathfrak{S}$ preserves limits of just the right type of ω^{op}-chains for our purposes.

3.5. *Coalgebra*

In Section 2 we have already encountered both non-deterministic and stochastic Kripke models, as well as the image finite variants of non-deterministic Kripke models. This plethora of different systems can be treated uniformly using coalgebras.

Definition 3.20. Let $\mathfrak{T}$ be a functor on a category **C**. A $\mathfrak{T}$-*coalgebra* is given by a **C**-object X together with a morphism $a : X \to \mathfrak{T}X$. We say a is the *dynamics* of the coalgebra. A *morphism* $f : (X, a) \to (Y, b)$ is given by a **C**-morphism $X \to Y$ such that

$$\begin{array}{ccc} X & \xrightarrow{\ f\ } & Y \\ {\scriptstyle a}\downarrow & & \downarrow{\scriptstyle b} \\ \mathfrak{T}X & \xrightarrow{\ \mathfrak{T}f\ } & \mathfrak{T}Y \end{array}$$

commutes.

Our primary examples of coalgebras will arise from the subprobability functor $\mathfrak{S}$: the coalgebras for $\mathfrak{S}$ are the stochastic relations $X \rightsquigarrow X$. Observe that we have to distinguish coalgebra morphisms as defined above from morphisms of stochastic relations proper as defined in Section 6.2. The latter are given by a pair of measurable functions instead of one, allowing us to model more general phenomena by treating the domain and codomain of the relation independently.

Other examples of coalgebras abound:

(1) Kripke frames arise as coalgebras for the powerset-functor $\mathfrak{P}$. Coalgebra morphisms for $\mathfrak{P}$ correspond to bounded morphisms as discussed in Section 2.2. Kripke models from a set Var of variables arise as coalgebras from the functor $X \mapsto \mathfrak{P}(X) \times \mathfrak{P}(\mathsf{Var})$.
(2) τ-Kripke frames for a modal similarity type $\tau = (O, \mathsf{ar})$ arise as coalgebras for the functor given by $X \mapsto \prod_{\Delta \in O} \mathfrak{P}(X^{\mathsf{ar}\Delta})$.
(3) Stochastic τ-Kripke frames for $\tau = (O, \mathsf{ar})$ arise as coalgebras for the functor given by $X \mapsto \prod_{\Delta \in O} \mathfrak{S}(X^{\mathsf{ar}\Delta})$. Here, $\mathfrak{S}$ is the subprobability functor.
(4) Image finite Kripke-frames arise as coalgebras for the *finite powerset functor* $\mathfrak{P}_\omega$, given by $\mathfrak{P}_\omega(X) = \{ A \subseteq X \mid A \text{ finite} \}$.

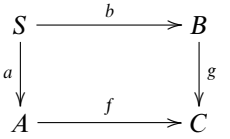
For a most influential early survey on coalgebras, see Rutten (2000). Coalgebras over measurable spaces are treated in depth in Doberkat (2007b).

Finally, we need the notion of a semi-pullback: this is a span, given a cospan, and is definded formally as follows.

Definition 3.21. Let $A \xrightarrow{\ f\ } C \xleftarrow{\ g\ } B$ be a cospan in a category **C**, that is, A, B, C are objects in **C** with morphisms f and g. A *semi-pullback* for the cospan is an object

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

S in $\mathbf{C}$ together with morphisms $A \xleftarrow{a} S \xrightarrow{b} B$ making the following diagram commutative:



Note that a semi-pullback only requires the object S together with the appropriate morphisms, and does not require a factorisation or a uniqueness property, as a pullback does.

4. Coalgebraic logic using predicate liftings

We now want to give generalised semantics for modal logics based on coalgebras over an arbitrary base category. In order to do this, two questions have to be settled:

(1) What are the admissible predicates over an abstract object?
(2) How are modal operators interpreted?

We give an answer to these two questions in Section 4.1. The notion of an admissible predicate is given axiomatically: for a chosen base category $\mathbf{C}$, we fix a functor $\Omega : \mathbf{C}^{\mathrm{op}} \to \mathbf{Set}$ and think of ΩC as the set of predicates over a $\mathbf{C}$-object C. This is akin to the development in modal logic, where the step from Kripke frames to so-called generalised frames, which are frames with a chosen family of subsets, has turned out to be very fruitful (Blackburn *et al.* 2001).

The second question is settled using (a slight reformulation) of the predicate liftings introduced by Pattinson (Pattinson 2004). Predicate liftings allow us to derive a predicate transformer $\Omega X \to \Omega X$ from a coalgebra (X, a). This predicate transformer is then used to give the semantics of a modal operator using the coalgebra (X, a). After defining the semantics of modal languages in coalgebras in Section 4.1, we present some examples of predicate liftings, both over the category **Set** and over the category **Meas** in Section 4.2. Section 4.3 sets out one of the main concerns of the current paper: the problem of expressivity. We argue that states x and y in coalgebras (X, a) and (Y, b), respectively, show the same observable behaviour if and only if there exists a coalgebra (Z, c) and coalgebra morphisms

$$(X, a) \xrightarrow{f} (Z, c) \xleftarrow{g} (Y, b)$$

with $f(x) = g(y)$. We say such states x and y are behaviourally equivalent. Thus, our slogan is 'observable behaviour is invariant under behavioural equivalence'. It is easy to see that behaviourally equivalent states satisfy exactly the same formulas. We are interested in the converse question: are states satisfying exactly the same formulas behaviourally equivalent? A logic with this property is said to be *expressive*.

In Section 4.4 we give a brief overview of a strategy for proving that a given coalgebraic logic is expressive. Finally, Section 4.5 recalls a result of Schröder on the expressivity of

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

coalgebraic logic over **Set**. In Section 5 we will present results on expressivity for functors over **Meas**.

4.1. *Motivation*

In the following, we let the category $\mathbf{C}$ and the functor Ω denote either:

— the category **Set** and the contravariant powerset functor $\mathcal{P}^- : \mathbf{Set}^{\mathrm{op}} \to \mathbf{Set}$; or
— the category **Meas** and the measurable subset functor $\mathcal{B} : \mathbf{Meas}^{\mathrm{op}} \to \mathbf{Set}$.

We think of ΩX as the set of admissible predicates over X. Observe that in each case ΩX is a Boolean algebra and Ωf acts through the inverse image, thus, $\Omega f(B) = f^{-1}[B]$. In particular, Ωf preserves countable meets and joins.

Take any functor $\mathfrak{T}$ on $\mathbf{C}$. We want to be able to define modal operators from $\mathfrak{T}$-coalgebras (X, a). To do this, we first seek to derive an interpretation of a (unary) modal operator as a predicate transformer on ΩX. Thus we seek to define a function $\Omega X \to \Omega X$ from the coalgebra's dynamics a. Recalling the definition of the Kripke semantics of the modal operator $\Diamond$ from Section 2.1, we rewrite this definition as follows:

$$x \models \Diamond\phi \iff a(x) \in \{\, B \in \mathcal{P}X \mid [\![\phi]\!] \cap B \neq \varnothing \,\}$$

where $a : X \to \mathcal{P}X$ is the successor relation of the Kripke model. Thinking of $\mathfrak{T}X$ as the 'shape' of possible successors assigned to x by the dynamics a, we thus *lift* a given predicate $[\![\phi]\!]$ to a predicate on $\mathfrak{T}X$ and ask whether $a(x)$ satisfies the lifted predicate. Now, let us write, for a set X and $A \subseteq X$,

$$\Diamond_X(A) = \{\, B \subseteq X \mid B \cap A \neq \varnothing \,\}.$$

So $\Diamond$ is a natural transformation $\mathcal{P}^- \to \mathcal{P}^- \cdot \mathcal{P}$. Observe that we can now write the extension of a formula in the following pointfree form:

$$[\![\Diamond\phi]\!]_{(X,a)} = \mathcal{P}^- a \cdot \Diamond_X([\![\phi]\!]_{(X,a)}). \tag{4}$$

This leads us to the following definition.

Definition 4.1. A *predicate lifting* for $\mathfrak{T} : \mathbf{C} \to \mathbf{C}$ is a natural transformation $\lambda : \Omega^n \to \Omega \cdot \mathfrak{T}$ for some natural number n. We say n is the *arity* of λ and write $\mathsf{ar}(\lambda) = n$.

The naturality of a predicate lifting $\lambda : \Omega^n \to \Omega \cdot \mathfrak{T}$ means, because of the contravariant nature of the functor involved, that the following diagram commutes in the category **Set** for each $\mathbf{C}$-morphism $f : X \to Y$:

$$\begin{array}{ccc} Y & (\Omega Y)^n & \xrightarrow{\lambda_Y} \Omega(\mathfrak{T}Y) \\ f \uparrow & {\scriptstyle (\Omega f)^n}\downarrow & \downarrow{\scriptstyle \Omega(\mathfrak{T}f)} \\ X & (\Omega X)^n & \xrightarrow{\lambda_X} \Omega(\mathfrak{T}X) \end{array}$$

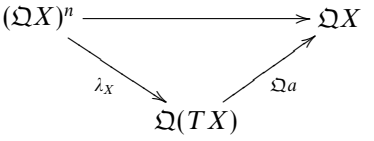
Thus, the following equation is required to hold for each $(B_i) \in (\Omega Y)^n$:

$$\lambda_X\big(\Omega f(B_1), \ldots, \Omega f(B_n)\big) = \Omega(\mathfrak{T}f)\big(\lambda_Y(B_1, \ldots, B_n)\big).$$

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

We will give some examples of predicate liftings in Section 4.2, but some general remarks are in order first.

We define from a predicate lifting λ and a coalgebra (X, a), a *predicate transformer* $(\mathfrak{Q}X)^n \to \mathfrak{Q}X$ as follows:



This will be used below to define the semantics of the modal language of some similarity type τ by using a glorified version of formula (4). First we need to consider a suitable notion of model in our scenario.

4.1.1. *Models.* Propositional connectives such as $\vee$, $\neg$, $\wedge$ are interpreted through the Boolean-algebraic structure of $\mathfrak{Q}X$. Modal operators, in turn, are interpreted through the coalgebra's dynamics a with the help of predicate liftings. We want to add the possibility of interpreting propositional variables. In the case of a coalgebraic modal logic over **Set**, this can be achieved by switching from the functor $\mathfrak{T}$ to $\mathfrak{T} \times \mathfrak{P}\mathsf{Var}$, where Var is some set of variables. Over **Meas**, we are not in a position to change the functors in this way due to measurability conditions. Instead, we have to add the possibility of interpreting propositional variables by switching to *models*.

Definition 4.2. Fix a set Var of propositional variables. A *model* is a $\mathfrak{T}$-coalgebra (X, a) together with a valuation $V : \mathsf{Var} \to \mathfrak{Q}X$. A *model morphism* $(X, a, V) \to (Y, b, W)$ is given by a coalgebra morphism $f : (X, a) \to (Y, b)$ such that $\mathfrak{Q}f \cdot W = V$ holds.

4.1.2. *Using predicate liftings to interpret modal logics.* Fix a functor $\mathfrak{T}$. In order to interpret the basic modal logic induced by a modal similarity type $\tau = (O, \mathsf{ar})$ as in Section 2.1, we associate an n-ary predicate lifting $[\![\Delta]\!] : \mathfrak{Q}^n \to \mathfrak{Q} \cdot \mathfrak{T}$ to each modal operator Δ of arity n. Having fixed such an association $[\![-]\!]$, we interpret the basic modal language based on τ and Var in any $\mathfrak{T}$-model $\mathsf{M} = (X, a, V)$ by assigning $[\![\phi]\!]_{\mathsf{M}} \in \mathfrak{Q}X$ to every formula ϕ as follows:

$$\begin{aligned}
[\![\top]\!]_{\mathsf{M}} &:= X \\
[\![\phi_1 \wedge \phi_2]\!]_{\mathsf{M}} &:= [\![\phi_1]\!]_{\mathsf{M}} \cap [\![\phi_2]\!]_{\mathsf{M}} \\
[\![\neg\phi]\!]_{\mathsf{M}} &:= X \setminus [\![\phi]\!]_{\mathsf{M}} \\
[\![v]\!]_{\mathsf{M}} &:= V(v) \ (v \in \mathsf{Var}) \\
[\![\Delta(\phi_1, \ldots, \phi_{\mathsf{ar}(\lambda)})]\!]_{\mathsf{M}} &:= \mathfrak{Q}a \cdot [\![\Delta]\!]_X([\![\phi_1]\!]_{\mathsf{M}}, \ldots, [\![\phi_{\mathsf{ar}(\lambda)}]\!]_{\mathsf{M}}).
\end{aligned} \tag{5}$$

For a state x of M, we write, as in Section 2.1,

$$\mathrm{Th}_{\mathsf{M}}(x) = \{\, \phi \in \mathscr{L}(\tau) \mid x \in [\![\phi]\!]_{\mathsf{M}} \,\}$$

and

$$\mathsf{M}, x \models \phi \iff x \in [\![\phi]\!]_{\mathsf{M}}.$$

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Definition 4.3. States x and y in models M and N are *logically equivalent* (with respect to a given modal similarity type τ and to interpretation $[\![-]\!]$) if and only if $\mathrm{Th}_{\mathsf{M}}(x) = \mathrm{Th}_{\mathsf{N}}(y)$ holds. The models M and N are *logically equivalent* if and only if

$$\{\, \mathrm{Th}_{\mathsf{M}}(x) \mid x \in X \,\} = \{\, \mathrm{Th}_{\mathsf{N}}(y) \mid y \in Y \,\}.$$

Thus, for every x we can find y with $\mathrm{Th}_{\mathsf{M}}(x) = \mathrm{Th}_{\mathsf{N}}(y)$, and *vice versa.*

4.2. *Examples of predicate liftings*

We are now going to give some examples of families of predicate liftings over **Set** or **Meas** and relate them to the modal logics discussed in Section 2.1.

4.2.1. *Predicate liftings for non-deterministic Kripke models.* We consider the category **Set** and the contravariant power set functor $\mathfrak{P}^-$ first. Examples of predicate liftings for endofunctors on **Set** include the following.

Examples 4.4.

(1) We have already seen in Section 4.1 that $[\![\Diamond]\!]$ given by

$$[\![\Diamond]\!]_X(A) = \{\, B \subseteq X \mid A \cap B \neq \varnothing \,\}$$

induces the usual semantics of the possibility modality $\Diamond$.

(2) The necessity modality $\Box$ can also be given a coalgebraic interpretation: use the predicate lifting $[\![\Box]\!] : \mathfrak{P}^- \to \mathfrak{P}^- \cdot \mathfrak{P}$ with

$$[\![\Box]\!]_X(A) = \{\, B \subseteq X \mid B \subseteq A \,\}.$$

Spelling out the coalgebraic semantics (5), we arrive at

$$x \models \Box\phi \iff \forall y \in a(x) : y \models \phi$$

for a $\mathfrak{P}$-coalgebra (X, a), and thus, at the usual Kripke semantics of $\Box$.

(3) For a modal similarity type $\tau = (O, \mathsf{ar})$, we can assign, to each $\Delta \in O$ of arity n, a natural transformation $[\![\Delta]\!] : \mathfrak{P}^{-n} \to \mathfrak{P}^- \cdot \mathfrak{P}(-)^n$ through

$$[\![\Delta]\!]_X(A_1, \ldots, A_n) = \{\, B \in \mathfrak{P}(X^n) \mid (A_1 \times \cdots \times A_n) \cap B \neq \varnothing \,\}.$$

This leads us to the usual Kripke semantics recalled in Section 2.1.

4.2.2. *Predicate liftings over measurable spaces.* We now consider **Meas** together with the measurable subsets functor $\mathfrak{B} : \mathbf{Meas}^{\mathrm{op}} \to \mathbf{Set}$. The following are examples of predicate liftings for various functors.

Examples 4.5.

(1) For the subprobability functor $\mathfrak{S}$, we have, for each real number r, a unary predicate lifting $\beta_r : \mathfrak{B} \to \mathfrak{B} \cdot \mathfrak{S}$ with

$$\beta_{r,X}(A) = \{\, \mu \in \mathfrak{S}(X) \mid \mu(A) \geqslant r \,\}. \tag{6}$$

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

We will now discuss β_r in some detail. Since $\beta_{r,X}(A) = \mathsf{ev}_A^{-1}[r,1]$, each $\beta_{r,X}(A)$ is a measurable subset of $\mathfrak{S}X$. Naturality of β_r boils down to the equation

$$\beta_{r,X}(f^{-1}[B]) = (\mathfrak{S}f)^{-1}[\beta_{r,Y}(B)]$$

for each measurable function $f : X \to Y$, $B \in \mathfrak{B}Y$, which in turn holds by the definition of $\mathfrak{S}$.

Spelling out the coalgebraic semantics for the modal operator $\Diamond_q$ induced by β_q, where q is any rational number, we obtain the semantics introduced in Section 2.1:

$$(X,a), s \models \Diamond_q \phi \iff a(s)([\![\phi]\!]) \geqslant q$$

for a stochastic relation (X,a). In summary, setting $[\![\Diamond_q]\!] = \beta_q$ allows us to treat the modal logic of stochastic relations coalgebraically. We will write

$$U = \{\, \beta_q \mid q \in \mathbb{Q} \cap [0,1] \,\}$$

for this set of predicate liftings.

(2) For a modal similarity type $\tau = (O, \mathsf{ar})$, we derive, for each $\Delta \in O$ of arity n and each rational number q, a predicate lifting

$$[\![\Delta_q]\!] : \mathfrak{B}^n \to \mathfrak{B}(\mathfrak{S}(-)^n)$$

through

$$[\![\Delta_q]\!]_X(A_1, \ldots, A_n) = \{\, \mu \in \mathfrak{S}(X^n) \mid \mu(A_1 \times \cdots \times A_n) \geqslant q \,\}.$$

The coalgebraic semantics induced by $[\![\Delta_q]\!]$ is the one presented in Section 3.3.1. Hence the semantics of the extended modal language can also be presented coalgebraically.

(3) Let Act be a countable set of actions. For each rational number q and each $a \in \mathsf{Act}$, we define a predicate lifting $\beta_{q,a}$ for $\mathfrak{S}^{\mathsf{Act}}$ through

$$\beta_{q,a,X}(B) = \{\, (\mu_i)_{i \in \mathsf{Act}} \mid \mu_a(B) \geqslant q \,\}.$$

We will write U^{Act} for this set of predicate liftings.

(4) For the functor $\mathfrak{S} \cdot \mathfrak{S}$, we may define a set $U \cdot U$ of predicate liftings as follows: the elements of $U \cdot U$ are indexed by finite sequences $(q_i) \in (\mathbb{Q} \cap [0,1])^{n+1}$ for $n \in \mathbb{N}$. For such a sequence of length $n+1$, we define an n-ary predicate lifting $\beta_{(q_i)}$ by

$$\begin{aligned}\beta_{(q_i),X}(B_1, \ldots, B_n) &= \{\, M \in \mathfrak{S}\mathfrak{S}X \mid M\big(\textstyle\bigcap_{i=1}^n \{\, \mu \in \mathfrak{S}X \mid \mu(B_i) \geqslant q_i \,\}\big) \geqslant q_{n+1} \,\} \\ &= \{\, M \in \mathfrak{S}\mathfrak{S}X \mid M\big(\textstyle\bigcap_{i=1}^n \beta_{q_i,X}(B_i)\big) \geqslant q_{n+1} \,\}.\end{aligned}$$

In Section 5.4 we will see that examples (3)–(4) are derived from example (1) generically. First we will discuss one of the main properties studied in the current paper: the problem of *expressivity*.

4.3. *The problem of expressivity*

Coalgebras, especially over **Set**, have been used to model a plethora of systems. For instance, coalgebras for the functor $\mathfrak{F}$ with $\mathfrak{F}X = A \times X$ model systems with output from the alphabet A. Now $\mathfrak{F}$ has a terminal coalgebra, given by the set A^ω of infinite streams

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

over the alphabet A. For any coalgebra morphism, the unique morphism into A^ω gives precisely the observable behaviour of the system. Note that states x and y in (possibly different) $\mathfrak{F}$-coalgebras show the same observable behaviour if and only if there exist morphism f and g with common codomain such that $f(x) = g(y)$. Abstracting away from this special case (and the existence of a terminal coalgebra), and, moreover, considering models, we arrive at the following definition.

Definition 4.6. States x and y of models M and N are said to be *behaviourally equivalent* if and only if there exists a diagram

$$\mathsf{M} \xrightarrow{f} \mathsf{L} \xleftarrow{g} \mathsf{N}$$

of model morphisms such that $f(x) = g(y)$. We say the models M and N are *behaviourally equivalent* if there exist *surjective* model morphisms f, g as above. The model L is said to be *co-mediating*.

Observe that this definition applies not only to models over the category **Set** but to models over arbitrary base categories of structured sets, such as **Meas**.

We have defined, using the conditions in (5), an interpretation of the logic induced by a modal similarity type τ in coalgebras by assigning a predicate lifting to each modal operator of τ. We take a slightly different viewpoint in this section, treating a set Λ of predicate liftings as primitive, and deriving from Λ a modal similarity type (Λ, ar) with ar as the arity-function of Λ. To distinguish a modal operator notationally from the predicate liftings, we write $\langle\lambda\rangle$ for the operator induced by λ. Of course, we always choose the original transformation as the interpretation. In symbols: $[\![\langle\lambda\rangle]\!] = \lambda$.

We write $\mathscr{L}(\Lambda)$ for the *negation free* logic induced by this set of predicate liftings Λ. Thus, we take as the modal similarity type $\tau = (\langle\Lambda\rangle, \mathsf{ar})$, with $\langle\Lambda\rangle = \{\, \langle\lambda\rangle \mid \lambda \in \Lambda \,\}$, and ar as the obvious function. We assume that the set of propositional variables is denoted by Var. We write $\mathscr{L}_\neg(\Lambda)$ for the logic induced by Λ, which includes negation.

Assume we have fixed some set Λ of $\mathfrak{T}$-predicate liftings. The following Proposition implies that behaviourally equivalent states in models satisfy exactly the same formulas.

Proposition 4.7. If $f : \mathsf{M} \to \mathsf{N}$ is a model morphism, we have $\mathfrak{Q}f([\![\phi]\!]_\mathsf{N}) = [\![\phi]\!]_\mathsf{M}$ for each $\phi \in \mathscr{L}(\Lambda)$.

Proof. We proceed by structural induction on a formula ϕ and use the naturality of the predicate liftings; *cf.* Doberkat and Schubert (2009), where the uni-modal case was treated. □

Corollary 4.8. Behavioural equivalence implies logical equivalence, for both individual states and models.

Thinking of $\mathscr{L}(\Lambda)$ as a specification language for systems modelled by $\mathfrak{T}$-coalgebras, we are interested in the converse implication: given states x and y in possibly different $\mathfrak{T}$-coalgebras that satisfy exactly the same formulas, are they behaviourally equivalent? In other words, if they do not show the same behaviour, there should by a formula $\phi \in \mathscr{L}(\Lambda)$

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

that is satisfied in one of the states, but not in the other. This is called the problem of *expressivity*.

Definition 4.9. A set Λ of $\mathfrak{T}$-predicate liftings is *expressive* if logically equivalent states in $\mathfrak{T}$-models are behaviourally equivalent.

In Section 5 we will survey some results on the expressivity of logics over the category of measurable spaces. To help in the appreciation of these developments, we will recall a result over **Set**, but first we will present some general results on predicate liftings.

4.4. *Some general results for coalgebraic logic*

In establishing the expressivity of a coalgebraic logic, we will rely on the following strategy:

(1) We define an equivalence relation from the logic, which identifies precisely those states that satisfy exactly the same formulas of the logic. Thus, we define the *logical equivalence relation* ρ_{M} on the state space X of a model M as follows:

$$\begin{aligned} x \,\rho_{\mathsf{M}}\, y &\iff \forall \phi \in \mathscr{L}(\Lambda) : (x \in [\![\phi]\!]_{\mathsf{M}} \iff y \in [\![\phi]\!]_{\mathsf{M}}) \\ &\iff \mathrm{Th}_{\mathsf{M}}(x) = \mathrm{Th}_{\mathsf{M}}(y). \end{aligned} \tag{7}$$

We have already used this relation in Lemma 2.7.

(2) Some properties of the family of predicate liftings defining the logic are explored in order to show that the equivalence relation is a congruence (to be defined below). This allows us to construct a model structure on X/ρ_{M} such that the projection $\eta_{\rho_{\mathsf{M}}} : X \to X/\rho_{\mathsf{M}}$ becomes a model morphism.

(3) The canonical projection $\eta_{\rho_{\mathsf{M}}}$ is used to construct a cospan as needed.

We will state necessary conditions over **Set** in Section 4.5. Section 5 is then devoted to deriving suitable conditions over **Meas**. Before following this programme, we will classify predicate liftings in both situations.

4.4.1. *Classifying predicate liftings.* We again use $\mathbf{C}$ and $\mathfrak{Q}$ to denote either **Set** and $\mathfrak{P}^-$ or **Meas** and $\mathfrak{B}$, respectively. Observe that in both cases there is a natural isomorphism

$$\mathfrak{Q} \cong \mathbf{C}(-, \mathbf{2}) \tag{8}$$

where $\mathbf{2}$ is either the set $\{0,1\}$ in the case of **Set** or the measurable space $(\{0,1\}, \mathfrak{P}\{0,1\})$. The following helpful result was established in Schröder (2008).

Lemma 4.10. *n*-ary predicate liftings for $\mathfrak{T} : \mathbf{C} \to \mathbf{C}$ are in one-to-one correspondence with $\mathfrak{Q}\mathfrak{T}(\mathbf{2}^n)$.

Proof. The isomorphism (8) allows us to apply Yoneda's Lemma (Mac Lane 1997, Proposition III.2.1). □

Hence, *n*-ary predicate liftings are in one-to-one correspondence with subsets of $\mathfrak{T}(\mathbf{2}^n)$ when $\mathbf{C} = \mathbf{Set}$ and with measurable subsets of $\mathfrak{T}(\mathbf{2}^n)$ when $\mathbf{C} = \mathbf{Meas}$.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

4.5. *Expressive coalgebraic logic over* **Set**

In this section we collect some results for coalgebraic logic over set in order to contrast them later with the situation over measurable spaces.

In order to obtain general expressivity results over **Set**, we have to use a slightly stronger logic. For a regular cardinal κ, we define $\mathscr{L}^{\kappa}_{\neg}(\Lambda)$ by the following grammar:

$$\phi ::= \langle\lambda\rangle\phi \mid \neg\phi \mid \bigwedge_{i\in I}\phi_i$$

where $\lambda \in \Lambda$ and the cardinality of $|I|$ is strictly less than κ.

Recall that a **Set**-endofunctor $\mathfrak{T}$ is said to be κ-*accessible* if, given an arbitrary set X and $t \in \mathfrak{T}X$, we find a subset $A \subseteq X$ of cardinality less than κ such that t is in the $\mathfrak{T}$-image of the inclusion $A \hookrightarrow X$. Thus, the action of $\mathfrak{T}$ on all sets is in some sense determined by the action of $\mathfrak{T}$ on sets of cardinality less than κ. The following result was proved in Schröder (2008).

Theorem 4.11 (Schröder). If $\mathfrak{T}$ is a κ-accessible endofunctor on **Set** and the set Λ is separating in the sense that for each set X and each $t \neq t' \in \mathfrak{T}X$, there exists $B \subseteq X$ such that $\lambda_X(B)$ contains exactly one of t, t', then $\mathscr{L}^{\kappa}_{\neg}(\Lambda)$ is expressive.

The result can be seen as a coalgebraic generalisation of the Hennessy–Milner Theorem (Theorem 2.8): the finite powerset functor is $\aleph_0$-accessible. Note that Theorem 4.11 deals with behavioural equivalence, whereas the Hennessy–Milner theorem 2.8 deals with the relationship of bisimilarity and logical equivalence. It is easy to see that the two concepts of behavioural equivalence and bisimilarity are equivalent over the category of sets, provided the functor preserves weak pullbacks. This is discussed at length in the survey Rutten (2000).

Over the category of measurable spaces, the situation is not so clear cut. In fact, the relationship between bisimilarity, behavioural equivalence and logical equivalence is one of the main topics of this paper. In the next section we will deal with the expressiveness of coalgebraic logics over **Meas**. Section 6 will deal with bisimilarity.

Before continuing, note that the cardinality restriction in Theorem 4.11 is necessary. It is well known that the Hennessy–Milner theorem fails for Kripke models in which we allow the successor sets of a state to have arbitrary cardinality. Somewhat surprisingly, for the subprobability functor, finite logical connectors suffice as well, in fact, we just need finite conjunctions. The technical reason for this lies in the π-λ-Theorem (Theorem 3.1). Finite conjunctions are no longer sufficient as soon as we leave the realm of subprobabilities (Desharnais *et al.* 2009).

5. Logical equivalence versus behavioural equivalence

We will now, following the outline given in Section 4.4, show that certain coalgebraic logics are expressive over **Meas**. We will do this by showing that the relation of logical equivalence is a congruence for a given coalgebra. In contrast to the situation over the category **Set**, it is not sufficient to know that some function exists since we also have to ensure measurability.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Our approach is inspired by Schröder's Theorem (Theorem 4.11). Before proving an analogue to this result, we will first introduce and discuss the notions of generating and separating sets of predicate liftings in Section 5.1. Separating sets of predicate liftings are akin to the **Set**-based notion introduced on the fly in the statement of Theorem 4.11. The notion of generating predicate lifting comes into play when showing that some constructed functions are measurable, and it has no counterpart in **Set**-based coalgebraic logic.

In Section 5.2 we prove that the logic induced by a generating and separating set of predicate liftings ensures that enough formulas exist so that we can define a model structure on the factor space after identifying logically indistinguishable states. This is used in Section 5.3 to show that the logic induced by a generating and separating set of predicate liftings is expressive.

Having established the usefulness of the notions of generating and separating set of predicate liftings, we will study closure properties in Section 5.4. In Section 5.5, we will show that we can construct terminal coalgebras for a fairly extensive range of functors on **Meas**. This is interesting in the context of the present paper since terminal coalgebras contain a unique state for each satisfied theory, that is, each set of formulas that arise as the theory of some state in some coalgebra.

Finally, we pull things together in Section 5.6 and present some examples of expressive coalgebraic logics.

5.1. *Generating and separating sets of predicate liftings*

In this section we fix an endofunctor $\mathfrak{T}$ on **Meas**. For a set Λ of $\mathfrak{T}$-predicate liftings and a generator $\mathscr{G}$ of some measurable space X, we write

$$\Lambda(\mathscr{G}) := \{\, \lambda_X(G_i) \mid \lambda \in \Lambda, (G_i) \in \mathscr{G}^{\mathsf{ar}(\lambda)} \,\} \subseteq \mathfrak{B}(\mathfrak{T}X).$$

Definition 5.1. A family Λ of $\mathfrak{T}$-predicate liftings:

— is said to be *separating* if and only if for each measurable space X and each π-generator $\mathscr{G}$ of $\mathfrak{B}X$, we have that $\Lambda(\mathscr{G})$ separates points in $\mathfrak{T}X$; and
— is said to be *generating* if and only if for each measurable space X and each π-generator $\mathscr{G}$ of X the set $\Lambda(\mathscr{G})$ generates $\mathfrak{B}(\mathfrak{T}X)$.

As a first reality check, we note the following example.

Example 5.2. The set $U = \{\, \beta_r \mid r \in \mathbb{Q} \cap [0,1] \,\}$ of predicate liftings for the subprobability functor from Example 4.5(1) is separating and generating. This is immediate from Lemma 3.8.

Remark 5.3. Take Λ, X and $\mathscr{G}$ as above, and put

$$\lambda_X^{\mathscr{G}}(t) := \{\, (G_i) \in \mathscr{G}^{\mathsf{ar}(\lambda)} \mid t \in \lambda_X(G_i) \,\}.$$

Then $\Lambda(\mathscr{G})$ separates points if and only if the family

$$\big(\lambda_X^{\mathscr{G}} : |\mathfrak{T}X| \to \mathfrak{P}(\mathscr{G}^{\mathsf{ar}(\lambda)})\big)_{\lambda\in\Lambda},$$

is jointly monic.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Proposition 5.4. A generating set of $\mathfrak{T}$-predicate liftings is separating if and only if each $\mathfrak{T}X$ is point-separating.

Proof. $\Lambda(\mathfrak{B}X)$ is a generator for $\mathfrak{B}(\mathfrak{T}X)$ by assumption. Thus, if Λ is separating, $\mathfrak{B}(\mathfrak{T}X)$ separates points, hence $\mathfrak{T}X$ is point-separating. On the other hand, if $\mathfrak{B}(\mathfrak{T}X) = \sigma(\Lambda(\mathscr{G}))$ separates points, $\Lambda(\mathscr{G})$ also does so by Lemma 3.4. □

In particular, a functor $\mathfrak{T}$ such that there exists a measurable space X with $\mathfrak{T}X$ not point-separating cannot have a set of predicate liftings that is both generating and separating. Moreover, we shall see in Section 5.4 that any constant functor $\mathfrak{C}_X$, where X is a non-point separating measurable space, has generating, non-separating sets of predicate liftings.

5.2. *The congruence theorem*

We now are going to show that we obtain a new model upon identifying logically equivalent states of a model, provided the set of predicate liftings defining the logic is separating and generating.

In this section, we again fix a **Meas**-endofunctor $\mathfrak{T}$ with a set Λ of predicate liftings.

Definition 5.5. For a model $\mathsf{M} = (X, d, V)$, we define a collection $\mathscr{G}_\mathsf{M}$ of subsets of X/ρ_M by

$$\mathscr{G}_\mathsf{M} := \{\, \eta_{\rho_\mathsf{M}}[[\![\phi]\!]_\mathsf{M}] \mid \phi \in \mathscr{L}(\Lambda) \,\}, \tag{9}$$

and define a σ-algebra $\mathscr{E}_\mathsf{M}$ on X/ρ_M through

$$\mathscr{E}_\mathsf{M} := \sigma(\mathscr{G}_\mathsf{M}),$$

where ρ_M is the equivalence relation (7) induced by the logic, and $\eta_{\rho_\mathsf{M}} : X \to X/\rho_\mathsf{M}$ is the projection.

Recall that $(x, y) \in \rho_\mathsf{M}$ holds if and only if $\mathrm{Th}_\mathsf{M}(x) = \mathrm{Th}_\mathsf{M}(y)$, and that every $[\![\phi]\!]_\mathsf{M}$ is ρ_M-invariant.

The σ-algebra $\mathscr{E}_\mathsf{M}$, which is induced by the extensions of formulas, was introduced and investigated in Doberkat (2009a).

Lemma 5.6. The following equation holds for every model $\mathsf{M} = (X, d, V)$:

$$\mathscr{G}_\mathsf{M} = \{\, A \subseteq X/\rho_\mathsf{M} \mid \eta_{\rho_\mathsf{M}}^{-1}[A] = [\![\phi]\!]_\mathsf{M} \text{ for some } \phi \in \mathscr{L}(\Lambda) \,\}.$$

Hence, $\eta_{\rho_\mathsf{M}} : X \to (X/\rho_\mathsf{M}, \mathscr{E}_\mathsf{M})$ is measurable. Moreover, $\mathscr{G}_\mathsf{M}$ is closed under finite intersections, and thus forms a π-generator of $\mathscr{E}_\mathsf{M}$.

Proof. First observe that each $[\![\phi]\!]_\mathsf{M}$ is ρ_M-invariant by construction. Hence, we have $\eta_{\rho_\mathsf{M}}^{-1}\left[\eta_{\rho_\mathsf{M}}[[\![\phi]\!]_\mathsf{M}]\right] = [\![\phi]\!]_\mathsf{M}$ by Lemma 3.7. This gives the inclusion from left to right. On the other hand, take any $B \subseteq X/\rho_\mathsf{M}$ with $\eta_{\rho_\mathsf{M}}^{-1}[B] = [\![\phi]\!]_\mathsf{M}$ for some $\phi \in \mathscr{L}(\Lambda)$. Then

$$\eta_{\rho_\mathsf{M}}[[\![\phi]\!]_\mathsf{M}] = \eta_{\rho_\mathsf{M}}\left[\eta_{\rho_\mathsf{M}}^{-1}[B]\right] = B$$

holds by the surjectivity of η_{ρ_M}. This yields $B \in \mathscr{G}_\mathsf{M}$.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Finally, observe that $\{ [\![\phi]\!]_{\mathsf{M}} \mid \phi \in \mathscr{L}(\Lambda) \}$ is closed under finite intersections since $\mathscr{L}(\Lambda)$ has finite conjunctions. Since $\eta_{\rho_{\mathsf{M}}}^{-1}$ preserves (finite) intersections, the closedness of $\mathscr{G}_{\mathsf{M}}$ under finite intersections is obvious. □

In general, $\mathscr{E}_{\mathsf{M}}$ will *not* be the final σ-algebra with respect to $\mathfrak{B}X$ and $\eta_{\rho_{\mathsf{M}}}$ – see Doberkat (2009b, Example 2.6.7). Thus, $\mathscr{E}_{\mathsf{M}}$ might be properly included in the final σ-algebra induced by the projection $\eta_{\rho_{\mathsf{M}}}$. Using $\mathscr{E}_{\mathsf{M}}$, rather than the final σ-algebra, is crucial for the constructions to be discussed now (Schubert 2009b).

Theorem 5.7. Let Λ be a separating and generating set of predicate liftings. For every model M, there exists a unique model structure $\mathsf{M}/\rho_{\mathsf{M}}$ over X/ρ_{M} such that $\eta_{\rho_{\mathsf{M}}} : \mathsf{M} \to \mathsf{M}/\rho_{\mathsf{M}}$ is a model morphism.

Proof. We write $\mathsf{M} = (X, d, V)$, η for $\eta_{\rho_{\mathsf{M}}}$, and Q for $(X/\rho_{\mathsf{M}}, \mathscr{E}_{\mathsf{M}})$.

(1) We will first show that $(x, y) \in \rho_{\mathsf{M}}$ implies $(x, y) \in \ker(\mathfrak{T}\eta \cdot d)$. Since

$$\mathscr{G} := \mathscr{G}_{\mathsf{M}} = \{ \eta[[\![\phi]\!]] \mid \phi \in \mathscr{L}(\Lambda) \}$$

is a π-generator for $\mathscr{E}_{\mathsf{M}}$ and Λ is assumed to be separating, it suffices to show that

$$(x, y) \in \ker\left(\lambda_Q^{\mathscr{G}} \cdot \mathfrak{T}\eta \cdot d\right)$$

holds for each $\lambda \in \Lambda$. We fix $\lambda \in \Lambda$, take any

$$(G_i) = (\eta[[\![\phi_i]\!]]) \in \mathscr{G}^{\mathsf{ar}(\lambda)}$$

and observe that we have

$$\begin{aligned} \mathfrak{B}d \cdot \mathfrak{B}\mathfrak{T}\eta \cdot \lambda_Q(G_i) &= \mathfrak{B}d \cdot \lambda_X \cdot \mathfrak{B}\eta^{\mathsf{ar}(\lambda)}(G_i) \\ &= \mathfrak{B}d \cdot \lambda_X(\eta^{-1}[\eta[[\![\phi_i]\!]]]) \\ &\overset{(\dagger)}{=} \mathfrak{B}d \cdot \lambda_X([[\![\phi_i]\!]]) \\ &= [[\![\langle\lambda\rangle(\phi_i)]\!]], \end{aligned} \tag{10}$$

where equation (†) holds since $[[\![\phi_i]\!]]$ is ρ_{M}-invariant. Using this fact, we obtain

$$\begin{aligned} (G_i) \in \lambda_Q^{\mathscr{G}} \cdot \mathfrak{T}\eta \cdot d(x) &\iff \mathfrak{T}\eta(d(x)) \in \lambda_Q(G_i) \\ &\iff x \in \mathfrak{B}d \cdot \mathfrak{B}\mathfrak{T}\eta \cdot \lambda_Q(G_i) \\ &\iff x \in [[\![\langle\lambda\rangle(\phi_i)]\!]]. \end{aligned}$$

Thus $x \,\rho_{\mathsf{M}}\, y$ implies

$$\lambda_Q^{\mathscr{G}} \cdot \mathfrak{T}\eta \cdot d(x) = \lambda_Q^{\mathscr{G}} \cdot \mathfrak{T}\eta \cdot d(y).$$

Hence $\rho_{\mathsf{M}} \subseteq \ker(\mathfrak{T}\eta \cdot d)$, which implies that there exists a unique function $q : Q \to \mathfrak{T}Q$ making

$$\begin{array}{ccc} X & \xrightarrow{\eta} & Q \\ {\scriptstyle d}\downarrow & & \downarrow{\scriptstyle q} \\ \mathfrak{T}X & \xrightarrow{\mathfrak{T}\eta} & \mathfrak{T}Q \end{array}$$

commute.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

(2) We need to show that q is measurable. For λ and (G_i) as above, we have

$$\mathfrak{B}\eta \cdot \mathfrak{B}q \cdot \lambda_Q(G_i) = \mathfrak{B}d \cdot \mathfrak{B}\mathfrak{T}\eta \cdot \lambda_Q(G_i) = [\![\langle\lambda\rangle(\phi_i)]\!]$$

by (10). Thus,

$$q^{-1}[\lambda_Q(G_i)] = \mathfrak{B}q \cdot \lambda_Q(G_i) \in \mathscr{G}$$

by Lemma 5.6. The measurability of q then follows since $\Lambda(\mathscr{G})$ forms a generator of $\mathfrak{B}\mathfrak{T}Q$ by assumption.

(3) We are left to define a valuation W on Q. For $v \in \mathsf{Var}$, we set

$$W(v) = \eta[\![v]\!]_{\mathsf{M}} \in \mathscr{E}_{\mathsf{M}}.$$

By ρ_{M}-invariance of $[\![v]\!]_{\mathsf{M}}$, we obtain

$$\mathfrak{B}\eta \cdot W(v) = \eta^{-1}[\eta[V(v)]] = V(v),$$

so W makes η a model morphism. The uniqueness of W with this property then holds by the injectivity of $\mathfrak{B}\eta$. □

5.2.1. *Models over analytic spaces.* When the underlying measurable space A of the model M is analytic (which is the blanket assumption in Doberkat (2007b), Doberkat and Schubert (2009) and Schubert (2007)), we obtain a well-behaved model structure on A/ρ_{M} even when Λ is not generating. Recall that a measurable space is said to be *analytic* if its measurable subsets arise as the Borel sets of a topological space that is the continuous image of a metrisable topological space with a countable base.

Theorem 5.8. Let Λ be a countable, separating set of predicate liftings, and let Var be countable. For every model $\mathsf{M} = (A, d, V)$ with A analytic, the σ-algebra $\mathscr{E}_{\mathsf{M}}$ is final with respect to the projection, and there exists a unique model structure $\mathsf{M}/\rho_{\mathsf{M}}$ over A/ρ_{M} such that $\eta_{\rho_{\mathsf{M}}} : \mathsf{M} \to \mathsf{M}/\rho_{\mathsf{M}}$ is a model morphism.

Proof. See Schubert (2009a). □

5.3. *Expressivity*

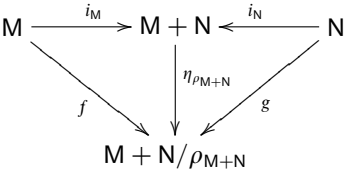
We will now show that the logic $\mathscr{L}(\Lambda)$ determined by a separating and generating set Λ of predicate liftings is *expressive*, that is, that the notions of logical and behavioural equivalence coincide. This entails that the logic is strong enough to distinguish states x and y which do not show the same behaviour, that is, for which we cannot find model morphisms f, g with $f(x) = g(y)$, *cf.* Section 4.3.

Theorem 5.9. Assume that Λ is separating and generating. If states x and y of models M and N, respectively, are logically equivalent, then they are behaviourally equivalent. Thus, the logic $\mathscr{L}(\Lambda)$ is expressive.

Proof. This is proved in the same way as in Schubert (2009a). We first observe that the category of models has (arbitrary) coproducts – see Doberkat and Schubert (2009). Thus,

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

given models M and N, we may form their coproduct $\mathsf{M}+\mathsf{N}$ in the category of models and then factor $\mathsf{M}+\mathsf{N}$ through $\rho_{\mathsf{M}+\mathsf{N}}$. This gives us the diagram



with i_{M}, i_{N} the coproduct injections and

$$f := \eta_{\rho_{\mathsf{M}+\mathsf{N}}} \cdot i_{\mathsf{M}}$$
$$g := \eta_{\rho_{\mathsf{M}+\mathsf{N}}} \cdot i_{\mathsf{N}}.$$

Logical equivalence of states x in M and y in N, respectively, implies

$$\mathrm{Th}_{\mathsf{M}+\mathsf{N}}(i_{\mathsf{N}}(y)) = \mathrm{Th}_{\mathsf{N}}(y) = \mathrm{Th}_{\mathsf{M}}(x) = \mathrm{Th}_{\mathsf{M}+\mathsf{N}}(i_{\mathsf{M}}(x)).$$

Thus, if x and y are logically equivalent, we have $f(x) = g(y)$. Hence $\mathsf{M}+\mathsf{N}/\rho_{\mathsf{M}+\mathsf{N}}$ is the required mediating model. □

The following result on the equivalence of the global properties is in fact an easy consequence of Theorem 5.9 – for a proof, see Schubert (2009a).

Theorem 5.10. Assume that Λ is separating and generating. Then logically equivalent models are behaviourally equivalent.

Example 5.11. The logic $\mathscr{L}(U)$ is expressive for $\mathfrak{S}$-coalgebras (thus taking $\mathsf{Var} = \varnothing$). Observe that $\mathscr{L}(U)$ has modal operators of the form $\Diamond_r$ for each $r \in \mathbb{Q} \cap [0,1]$, which are interpreted in a $\mathfrak{S}$-coalgebra (X, a) through

$$x \in [\![\Diamond_r \phi]\!] \iff a(x)([\![\phi]\!]) \geqslant r.$$

The expressivity of this logic (and its generalisation using U^{Act} to functors of the form $\mathfrak{S}^{\mathsf{Act}}$ for some set Act of actions) was first shown in Danos *et al.* (2006) using a quite different notion of morphism, and in Jacobs and Sokolova (2010) and Doberkat (2009a) using coalgebraic methods.

Making use of Theorem 5.8, we may show that a countable logic interpreted using a separating set of predicate liftings is expressive *over analytic spaces*. The following corollary expresses this in detail.

Corollary 5.12. Assume that Λ is separating and that Λ and Var are countable. Let M and N be models whose underlying spaces are analytic. Then states x and y of M and N, respectively, are behaviourally equivalent if and only if they are logically equivalent. Moreover, the models M and N are behaviourally equivalent if and only if they are logically equivalent.

Thus, over analytic spaces, we may replace the requirement that the set of predicate liftings is generating by the requirement that it is countable.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

5.4. *Closure properties of predicate liftings*

We are now going to establish the fact that the class of functors on **Meas** for which a generating and separating class of predicate liftings exists is closed under limits, coproducts and composition. This enables us to construct expressive modal logics for endofunctors on **Meas** in a modular way. Similar modular construction have been studied in Cîrstea and Pattinson (2004), Cîrstea (2006) and Schröder (2008) for functors over **Set**.

5.4.1. *Constant functors.* Let A be a measurable space. The *constant functor* $\mathfrak{C}_A$ with value A maps every measurable space to A and every measurable function to the identity 1_A. Lemma 4.10 implies that unary predicate liftings for $\mathfrak{C}_A$ are in one-to-one correspondence with the measurable subsets of A: given $C \in \mathfrak{B}A$, we define a $\mathfrak{C}_A$-predicate lifting $c_C : \mathfrak{B} \to \mathfrak{B} \cdot \mathfrak{C}_A$ by $c_{C,X}(B) = C$ for X any measurable space and $B \in \mathfrak{B}(X)$. It is easy to see that every $\mathfrak{C}_A$-predicate lifting arises in this way.

Observe that a set Λ of predicate liftings for $\mathfrak{C}_A$ is:

— generating if and only if the corresponding subset of $\mathfrak{B}A$ is a generator;
— separating if and only if the corresponding subset of $\mathfrak{B}A$ separates points.

Thus every non-point-separating measurable space A gives an example of a functor that admits a generating set of predicate liftings that is not separating.

5.4.2. *Closure under limits.* Let $\mathfrak{T}$, $\mathfrak{T}_i$ be functors on **Meas** ($i \in I$) and suppose that each $\mathfrak{T}_i$ has a family Λ_i of predicate liftings. If $(\alpha_i : \mathfrak{T} \to \mathfrak{T}_i)_{i\in I}$ is a family of natural transformations, we obtain a family Λ of $\mathfrak{T}$-predicate liftings through

$$\Lambda = \{\, \mathfrak{B}\alpha_i \cdot \lambda \mid i \in I, \lambda \in \Lambda_i \,\}.$$

We say Λ is the *lift of the* Λ_i *along* $(\alpha_i)_{i\in I}$.

Proposition 5.13. If all Λ_i are generating, and the family $(\alpha_{i,X} : \mathfrak{T}X \to \mathfrak{T}_iX)_{i\in I}$ is initial for each measurable space X, then the lift Λ of the Λ_i along $(\alpha_i)_{i\in I}$ is generating as well.

Proof. Let X be a measurable space and $\mathscr{G}$ be a π-generator of X. By assumption, $\Lambda_i(\mathscr{G})$ is a generator for $\mathfrak{B}(\mathfrak{T}_iX)$ for each $i \in I$. Hence, by Lemma 3.3, $\bigcup_{i\in I} \alpha_{i,X}^{-1}[\Lambda_i(G)]$ is a generator for $\mathfrak{B}(\mathfrak{T}X)$. But this set is just $\Lambda(\mathscr{G})$. □

Proposition 5.14. If all Λ_i are separating, and the family $(\alpha_{i,X} : \mathfrak{T}X \to \mathfrak{T}_iX)_{i\in I}$ is jointly monic for each measurable space X, then Λ is separating too.

Proof. Let X be a measurable set with π-generator $\mathscr{G}$. For $t \neq t' \in \mathfrak{T}X$, we find $j \in I$ with $\alpha_{j,X}(t) \neq \alpha_{j,X}(t')$. Since Λ_j is assumed to be separating, we find $\lambda \in \Lambda_j$ and $(G_i) \in \mathscr{G}^{\mathsf{ar}\lambda}$ with, say, $\alpha_{j,X} \in \lambda_X(G_i)$ and $\alpha_{j,X}(t') \notin \lambda_X(G_i)$. Thus, $t \in \mathfrak{B}\alpha_{j,X} \cdot \lambda_X(G_i)$ and $t' \notin \mathfrak{B}\alpha_{j,X} \cdot \lambda_X(G_i)$, that is, Λ is separating. □

We will now use the previous two results to show that the class of functors that admit a separating and generating set of predicate liftings is closed under limits in a sense that we will now make precise. Recall, for instance from Mac Lane (1997, Corollary V.4), that the category whose objects are functors **Meas** → **Meas** and whose morphisms are natural

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

transformations has all (set-indexed) limits and that these limits are calculated pointwise; *cf.* Section 3.2. Thus, for example, the product $\prod_{i\in I}\mathfrak{T}_i$ of a family $(\mathfrak{T}_i : \mathbf{Meas} \to \mathbf{Meas})_{i\in I}$ of functors is formed through

$$(\prod_{i\in I}\mathfrak{T}_i)(X) = \prod_{i\in I}\mathfrak{T}_i(X),$$

and the latter product is taken in **Meas**.

Combining Propositions 5.13 and 5.14, we get that the class of functors for which a generating and separating set of predicate liftings exists is closed under limits.

Proposition 5.15. The lift of a family of generating and separating sets of predicate liftings along a limit source is generating and separating.

Proof. Let $(\ell_i : \mathfrak{T} \to \mathfrak{T}_i)_{i\in I}$ be a limit, that is, for each measurable space X, we have

$$(\ell_{i,X} : \mathfrak{T}X \to \mathfrak{T}_iX)_{i\in I}$$

is a limit source in **Meas**, and hence is initial and jointly monic. The result now follows from Propositions 5.13 and 5.14 □

We will demonstrate this construction by calculating the Act-fold power V of the set U of predicate liftings for the subprobability functor $\mathfrak{S}$. This is going to be a set of predicate liftings for $\mathfrak{S}^{\mathsf{Act}}$. We write $\pi_a : \mathfrak{S}^{\mathsf{Act}} \to \mathfrak{S}$ for the ath projection, $a \in \mathsf{Act}$. An element of V has the form $\mathfrak{B}\pi_a \cdot \beta_r$, for $a \in \mathsf{Act}$, $\beta_r \in U$, and we have

$$(\mu_i)_{i\in\mathsf{Act}} \in (\mathfrak{B}\pi_a \cdot \beta_r)_X(A) \iff \pi_{a,X}(\mu_i) \in \beta_{r,X}(A) \iff \mu_a(A) \geqslant r. \tag{11}$$

Thus, the set V is just the U^{Act} from Example 4.5(3). From this, we get that U^{Act} is separating and generating, and hence the corresponding logic $\mathscr{L}(U^{\mathsf{Act}})$ is expressive. This result has already been observed a couple of times in the literature (Danos *et al.* 2006; Doberkat 2009a).

5.4.3. *Closure under coproducts.* We will now show that the class of endofunctors on **Meas** that admit a separating and generating set of predicate liftings is also closed under coproducts. Sadly, the construction is slightly more involved than is the case for limits. Before we give the construction, we will make a few observations on coproducts of functors and their predicate liftings.

As with limits, the coproduct of a family $(\mathfrak{T}_i)_{i\in I}$ of **Meas**-endofunctors is calculated pointwise. Thus, writing $\coprod_{i\in I}\mathfrak{T}_i$ for the coproduct, we have

$$(\coprod_{i\in I}\mathfrak{T}_i)(X) = \coprod_{i\in I}\mathfrak{T}_iX,$$

and the latter coproduct is taken in **Meas** according to Section 3.2.

It turns out that predicate liftings for $\coprod_{i\in I}\mathfrak{T}_i$ are easy to describe in terms of families of predicate liftings for the individual $\mathfrak{T}_i$. Since $\mathfrak{B} : \mathbf{Meas}^{\mathrm{op}} \to \mathbf{Set}$ preserves limits, n-ary predicate liftings λ for $\coprod_{i\in I}\mathfrak{T}_i$ correspond to I-indexed families of n-ary predicate liftings $(\lambda_i : \mathfrak{B}^n \to \mathfrak{B}\cdot\mathfrak{T}_i)_{i\in I}$; we write $\lambda = [\lambda_i]$ for the $\coprod_{i\in I}\mathfrak{T}_i$-predicate lifting corresponding to the family $(\lambda_i)_{i\in I}$ of $\mathfrak{T}_i$-predicate liftings. We will use this observation to show that, given

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

separating and generating sets Λ_i of $\mathfrak{T}_i$-predicate liftings for each $i \in I$, we can find a separating and generating set of predicate liftings for $\coprod \mathfrak{T}_i$.

Definition 5.16. Let n be a natural number and $\mathfrak{T} : \mathbf{Meas} \to \mathbf{Meas}$ be a functor.

— We write $\bot_n^{\mathfrak{T}}$ for the constant n-ary predicate lifting with value $\varnothing$.
— We write $\top_n^{\mathfrak{T}}$ for the constant n-ary predicate lifting with value $\mathfrak{T}X$.
— For a set Λ of $\mathfrak{T}$-predicate liftings, we write $\Lambda|_n = \{\, \lambda \in \Lambda \mid \mathsf{ar}(\lambda) = n \,\}$.

Proposition 5.17. Let $(\mathfrak{T}_i)_{i\in I}$ be a set-indexed family of **Meas**-endofunctors such that each $\mathfrak{T}_i$ has a separating and generating set Λ_i of predicate liftings, and write $N := \mathsf{ar}[\bigcup_{i\in I} \Lambda_i] \subseteq \mathbb{N}$. Then

$$\Lambda = \bigcup_{n\in N} \{\, [\lambda_i] \mid (\lambda_i)_{i\in I} \in \prod_I (\Lambda_i|_n \cup \{\, \bot_n^{\mathfrak{T}_i}, \top_n^{\mathfrak{T}_i} \,\}) \,\}$$

is a separating and generating set of predicate liftings for $\mathfrak{T} = \coprod_{i\in I} \mathfrak{T}_i$.

Proof. Fix a measurable space X with π-generator $\mathscr{G}$. Take any $t, t' \in \mathfrak{T}X$ with $t \neq t'$. When t stems from $\mathfrak{T}_i X$ and t' stems from $\mathfrak{T}_j X$ with $i \neq j$, we can separate them using the unary predicate lifting $[\lambda_k]_{k\in I}$ with $\lambda_i = \bot_1^{\mathfrak{T}_i}$ and $\lambda_k = \top_1^{\mathfrak{T}_k}$ for all $k \neq i$.

When $t, t' \in \mathfrak{T}_i X$, we find $\lambda \in \Lambda_i$, $(G_k) \in \mathscr{G}^{\mathsf{ar}(\lambda)}$ with, say, $t \in \lambda_X(G_k)$, $t' \notin \lambda_X(G_k)$. We now define a family $(\kappa_j : \mathfrak{B} \to \mathfrak{B}\cdot\mathfrak{T}_j)_{j\in I}$ of predicate liftings by $\kappa_i = \lambda$ and $\kappa_j = \top_{\mathsf{ar}\lambda}^{\mathfrak{T}_j}$ for all $j \neq i$. Then the set $[\kappa_j](G_k) \in \mathfrak{B}(\mathfrak{T}X)$ contains t but does not contain t'.

Showing that Λ is generating provided each Λ_i is generating is straightforward. □

5.4.4. *Predicate liftings for composite functors.* Let $\mathfrak{R}$ and $\mathfrak{T}$ be endofunctors on **Meas** with sets K and Λ, respectively, of predicate liftings. We define a set $\Lambda * K$ of predicate liftings for $\mathfrak{T}\cdot\mathfrak{R}$ as follows (*cf.* Schröder (2008)):

$$\Lambda * K = \left\{\, \lambda\mathfrak{R} \cdot \textstyle\prod_{1,\dots,\mathsf{ar}(\lambda)} \kappa_i \mid \lambda \in \Lambda, (\kappa_i) \in K^{\mathsf{ar}(\lambda)} \,\right\}.$$

Thus, an element of $\alpha \in \Lambda * K$ can be presented by a pair $(\lambda, (\kappa_i))$ with $\lambda \in \Lambda$ and (κ_i) an $\mathsf{ar}(\lambda)$-indexed family of predicate liftings in K. We have $\mathsf{ar}(\alpha) = \sum_i \mathsf{ar}(\kappa_i)$ and

$$\alpha_X(A_1, \dots, A_{\mathsf{ar}(\lambda)}) = \lambda_{\mathfrak{R}X}\left(\kappa_{1,X}(A_1), \dots, \kappa_{\mathsf{ar}(\lambda),X}(A_{\mathsf{ar}(\lambda)})\right) \in \mathfrak{B}(\mathfrak{T}\cdot\mathfrak{R}(X)) \tag{12}$$

for each $(A_i)_{i\in\{1,\dots,\mathsf{ar}(\lambda)\}}$ with $A_j \in (\mathfrak{B}X)^{\mathsf{ar}(\kappa_j)}$ and X a measurable space.

Having now defined what could be a suitable notion of the composite of predicate liftings, we would like to show that $\Lambda * K$ is generating provided Λ and K are. At this point we face a difficulty: given a π-generator $\mathscr{G}$ of some measurable set X, we might conclude that $K(\mathscr{G})$ is a generator, but it will in general not be a π-generator, that is, closure under finite intersections is lost. For an example of this behaviour consider the set U of predicate liftings for the subprobability functor. This leads us to the following definition.

Definition 5.18. We say a set Λ of predicate liftings is *π-generating* if $\Lambda(\mathscr{G})$ is a π-generator for each π-generator $\mathscr{G}$.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

As we mentioned earlier, the set U from Examples 4.5(1) is not π-generating. However, the notion of a π-generating predicate lifting allows us to state the following result.

Lemma 5.19. The following statements hold for endofunctors $\mathfrak{R}$ and $\mathfrak{T}$ on **Meas** with sets K and Λ of predicate liftings, respectively:

(1) If K is π-generating and Λ is generating, then $\Lambda * K$ is generating.
(2) If K is π-generating and Λ is separating, then $\Lambda * K$ is separating.

Proof. We just need to observe that we have

$$(\Lambda * K)(\mathscr{G}) = \Lambda(K(\mathscr{G}))$$

for each π-generator $\mathscr{G}$. □

From each generating set Λ, we can derive a π-generating set by forming the π-*closure* $\pi(\Lambda)$; see below. Forming $\pi(\Lambda)$ is akin to forming the Boolean closure as in Schröder (2008).

In order to define $\pi(\Lambda)$, we first need to introduce some notation. For a finite set $F = \{1,\dots,n\}$ and an F-indexed family $(\lambda_i)_F$ of $\mathfrak{T}$-predicate liftings, we define a $\mathfrak{T}$-predicate lifting $\bigotimes_F \lambda_i$ of arity $\mathsf{ar}\big(\bigotimes_F \lambda_i\big) = \sum_F \mathsf{ar}(\lambda_i)$ through

$$\big(\bigotimes_F \lambda_i\big)_X(A_{1,1},\dots A_{1,\mathsf{ar}\lambda_1}, A_{2,1},\dots,A_{2,\mathsf{ar}\lambda_2},\ \dots\ ,A_{n,1},\dots,A_{n,\mathsf{ar}\lambda_n}) = \bigcap_{i\in F} \lambda_{i,X}(A_i),$$

with $A_{i,k} \in \mathfrak{B}X$, $i \in F$, $k \in \{1,\dots,\mathsf{ar}(\lambda_i)\}$. Naturality of $\bigotimes_F \lambda_i$ is obvious. When $F = \varnothing$, we have $\bigotimes_\varnothing = \top_0^{\mathfrak{T}}$ is the constant predicate lifting with value $\mathfrak{T}X \in \mathfrak{B}\mathfrak{T}X$. Finally, we define the π-*closure* of a family Λ of $\mathfrak{T}$-predicate liftings to be the following family of predicate liftings:

$$\pi(\Lambda) = \Big\{ \bigotimes_F \lambda_i \mid F = \{1,\dots,n\} \text{ for some } n \in \mathbb{N}, (\lambda_i) \in \Lambda^F \Big\}.$$

Observe that $\pi(\Lambda)$ contains predicate liftings of arbitrary arity even when Λ consists of unitary predicate liftings.

Lemma 5.20. The following statements are equivalent for a set Λ of $\mathfrak{T}$-predicate liftings:

(1) Λ is generating.
(2) $\pi(\Lambda)$ is π-generating.

Proof. For a π-generator $\mathscr{G}$ of some measurable space X, we have $\pi(\Lambda)(\mathscr{G}) \subseteq \sigma(\Lambda(\mathscr{G}))$ and $\Lambda(\mathscr{G}) \subseteq \pi(\Lambda)(\mathscr{G})$. Hence,

$$\sigma(\Lambda(\mathscr{G})) \subseteq \sigma(\pi(\Lambda)(\mathscr{G})) \subseteq \sigma(\sigma(\Lambda(\mathscr{G}))) = \sigma(\Lambda(\mathscr{G})).$$

Thus $\Lambda(\mathscr{G})$ generates $\mathfrak{B}\mathfrak{T}X$ if and only if $\pi(\Lambda)(\mathscr{G})$ generates $\mathfrak{B}\mathfrak{T}X$. In any case, $\pi(\Lambda)(\mathscr{G})$ is closed under finite intersections by construction. □

Theorem 5.21. Let $\mathfrak{R}$ and $\mathfrak{T}$ be endofunctors on **Meas**, K be a set of $\mathfrak{R}$-predicate liftings and Λ be a set of $\mathfrak{T}$-predicate liftings. If K is π-generating and Λ is generating and separating, then the logic $\mathscr{L}(\Lambda * K)$ is expressive.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Proof. $\Lambda * K$ is a generating and separating set of predicate liftings for $\mathfrak{T} \cdot \mathfrak{R}$ by Lemma 5.19, so the result follows from Theorem 5.9. □

Using the π-closure construction, we can get rid of the (quite strict) assumption that K is π-generating.

Corollary 5.22. Let $\mathfrak{R}$ and $\mathfrak{T}$ be endofunctors on **Meas**, K be a set of $\mathfrak{R}$-predicate liftings and Λ be a set of $\mathfrak{T}$-predicate liftings. If K is generating and Λ is generating and separating, the logic $\mathscr{L}(\Lambda * \pi(K))$ is expressive.

Observe that the logic $\mathscr{L}(\Lambda * \pi(K))$ has modal operators of arbitrary arity, even if K and Λ just consist of unitary predicate liftings.

Example 5.23. The set $U \cdot U$ introduced in Examples 4.5(4) is just $U * \pi(U)$. Thus, the composite $\mathfrak{S} \cdot \mathfrak{S}$ admits an expressive set of predicate liftings.

In order to summarise the closure properties derived so far, we make the following definition.

Definition 5.24. The class of *measure-polynomial endofunctors* on **Meas** is the smallest class of functors satisfying the following properties:

— It contains the identity and the subprobability functor.
— It contains every constant functor $\mathfrak{C}_X$, with X a measurable space.
— It is closed under composition.
— It is closed under limits and coproducts.

Theorem 5.25. Every measure-polynomial endofunctor $\mathfrak{T}$ with $\mathfrak{T}X$ point-separating for each X admits an expressive coalgebraic logic.

Observe that a measure-polynomial endofunctor satisfies the condition of Theorem 5.25 provided all the constant functors used in its definition take values in point-separating measurable spaces.

Proof. By Proposition 5.4, we just need to find a generating set Λ of predicate liftings in order to apply Theorem 5.9. We can define such a set by induction on the structure of $\mathfrak{T}$. For the identity, take the identity predicate lifting $1_{\mathfrak{B}}$. For the subprobability functor, take $\Lambda = U$ as in Example 4.5(1). For a constant functor $\mathfrak{C}_A$, let Λ be the set corresponding to $\mathfrak{B}A$. When $\mathfrak{T}$ arises by composition, say $\mathfrak{T} = \mathfrak{T}_2 \cdot \mathfrak{T}_1$, we assume that we have already fixed generating sets Λ_i for $\mathfrak{T}_i$, $i = 1, 2$. Then $\Lambda_2 * \pi(\Lambda_1)$ is a generating set of predicate liftings for $\mathfrak{T}$ by Lemma 5.19. The case where $\mathfrak{T}$ arises as a limit or as a coproduct can be treated using Proposition 5.15 or Proposition 5.17. □

The requirement in Theorem 5.25 that each $\mathfrak{T}X$ is point-separating is necessary, as the following counter example shows.

Example 5.26. Let A be a measurable space and consider the functors $\mathfrak{T}_1 : X \longmapsto \mathfrak{S}(X \times A)$ and $\mathfrak{T}_2 : X \longmapsto \mathfrak{S}(X) \times A$. Both can be used to model stochastic systems with output in the set A. Observe that $\mathfrak{T}_1$ admits an expressive coalgebraic logic, given by $U * \pi(\{1\} \times \mathfrak{B}A)$,

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

where we have written $\mathfrak{B}A$ for the set of constant predicate liftings corresponding to the measurable subsets of A.

On the other hand, $\mathfrak{T}_2$ fails to admit an expressive logic if A is not point-separating: take a, a' in A that cannot be separated by a measurable subset in A and define a coalgebra structure d on $\mathbf{2}$ by $d(0) = (\delta_0, a)$, $d(1) = (\delta_0, a')$ where δ_x is the Dirac-measure defined by

$$\delta_x(B) = \begin{cases} 1 & x \in B, \\ 0 & x \notin B. \end{cases}$$

Since coalgebraic logics, in our setting, allow only measurable subsets to arise as extensions of formulas, there cannot be any coalgebraic logic for which the states 1 and 0 are not logically equivalent. On the other hand, they cannot be behaviourally equivalent: given a coalgebra (Y, b) and morphisms $f, g : (\mathbf{2}, d) \to (Y, b)$ with $f(0) = g(1)$, we would have, writing $p_2 : \mathfrak{S}\mathbf{2} \times A \to A$ for the second projection,

$$a = p_2 \cdot \mathfrak{T}_2 f \cdot d(0) = p_2 \cdot b \cdot f(0) = p_2 \cdot b \cdot g(1) = p_2 \cdot \mathfrak{T}_2 g \cdot d(1) = a', \tag{13}$$

and hence a contradiction.

5.5. *Terminal coalgebras*

We will now shift our attention to the question as to whether terminal coalgebras exist. Recall that an object T in a category $\mathbf{C}$ is said to be *terminal* if for each $\mathbf{C}$-object C there is a unique morphism $!_C : C \to T$.

This seems to be quite remote from coalgebraic modal logic, but we will see that there is in fact an intimate connection between terminal coalgebras and coalgebraic logic.

The following definition is a slight modification of the notion of a measure-polynomial endofunctor – see Definition 5.24.

Definition 5.27. The class of *Standard Borel measure-polynomial functors* is the smallest class of functors that contains the identity, the subprobability functor and every constant functor $\mathfrak{C}_A$ with A Standard Borel, and is closed under composition and *countable* limits and coproducts.

Note that the image of a Standard Borel space under a Standard Borel measure-polynomial functor is again a Standard Borel space.

We start with the following theorem from Schubert (2009c):

Theorem 5.28. Let $\mathfrak{T}$ be a Standard Borel measure-polynomial endofunctor. Then there exists a terminal $\mathfrak{T}$-coalgebra whose underlying space is Standard Borel.

Proof. We will just sketch the proof given in Schubert (2009c). The *terminal sequence* of a Standard Borel measure-polynomial endofunctor $\mathfrak{T}$ defined as the following diagram:

$$1 \xleftarrow{!} \mathfrak{T}1 \xleftarrow{\mathfrak{T}!} \mathfrak{T}^2 1 \cdots \mathfrak{T}^n 1 \xleftarrow{\mathfrak{T}^n !} \mathfrak{T}^{n+1} 1 \cdots \qquad (n \in \mathbb{N}),$$

with 1 a one-point measurable space and $! : \mathfrak{T}1 \to 1$ the unique function. Observe that when $\mathfrak{T}1 = \varnothing$ we have $\mathfrak{T}X = \varnothing$ for all $X \in \mathbf{Meas}$, and hence $\varnothing$ is the only measurable

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

space on which a $\mathfrak{T}$-coalgebra-structure exists. Thus, we may safely assume that $\mathfrak{T}1 \neq \varnothing$ holds and thus that $! : \mathfrak{T}1 \to 1$ is a retraction. This in turn implies that each $\mathfrak{T}^n!$ is surjective.

We write $(L, (\ell_n)_{n\in\mathbb{N}})$ for the limit of this diagram. Using induction on the structure of $\mathfrak{T}$, we can show that $\mathfrak{T}$ preserves this limit. The only non-trivial part being the case $\mathfrak{T} = \mathfrak{S}$, which is taken care of by the Kolmogoroff Consistency Theorem 3.18.

Thus, we have an isomorphism $h : L \to \mathfrak{T}L$ in **Meas**, and (L, h) is a terminal $\mathfrak{T}$-coalgebra. So L is a Standard Borel space since the class of Standard Borel spaces is closed under countable limits in **Meas**. □

We fix a functor $\mathfrak{T}$ on **Meas** together with a generating and separating set Λ of predicate liftings for $\mathfrak{T}$ and assume that $\mathfrak{T}$ has a terminal coalgebra (L, h).

Proposition 5.29.

(1) States x and y of coalgebras (X, a) and (Y, b), respectively, are logically equivalent if and only if they are behaviourally equivalent if and only if $!_{(X,a)}(x) = !_{(Y,b)}(y)$ holds.
(2) States x and y of (L, h) are logically equivalent if and only if they are equal.

Proof.

(1) To prove the first statement, we only need to show that $f(x) = g(y)$ for some coalgebra morphisms f, g with common codomain (Z, c) implies $!_{(X,a)}(x) = !_{(Y,b)}(y)$. This follows from the uniqueness property of morphisms $(Z, c) \to (L, h)$: we have

$$\begin{aligned} !_{(X,a)}(x) &= !_{(Z,c)} \cdot f(x) \\ &= !_{(Z,c)} \cdot g(y) \\ &= !_{(Y,b)}(y). \end{aligned}$$

(2) The second statement follows since $!_{(L,h)}$ is the identity. □

Thus, the states of the terminal coalgebra (L, h) are in one-to-one correspondence with *satisfied theories*, that is, subsets of $\mathscr{L}(\Lambda)$ of the form $\mathrm{Th}_{(X,a)}(x)$ for some coalgebra (X, a), $x \in X$. This insight was used in Moss and Viglizzo (2006) to construct terminal coalgebras.

5.6. *Examples*

We will now give some examples of expressive coalgebraic logics for coalgebras over the category **Meas**. The results from Section 5.4 will be used to show that many of the measure-polynomial endofunctors (see Definition 5.24) in fact admit an expressive coalgebraic logic. This generalises several results from the literature.

5.6.1. *Probabilities.* We say $\mu \in \mathfrak{S}X$ is a *probability measure* on a measurable space X if $\mu(X) = 1$. We write $\mathfrak{S}_1$ for the subfunctor of $\mathfrak{S}$ given by the probability measures. Observe that we may and do restrict the set U of predicate liftings to $\mathfrak{S}_1$. Thus, we obtain the following logic for $\mathfrak{S}_1$-coalgebras (taking $\mathsf{Var} = \varnothing$):

$$\phi ::= \top \mid \phi_1 \wedge \phi_2 \mid \Diamond_q \phi$$

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

where q is a rational number in $[0,1]$. The interpretation in a $\mathfrak{S}_1$-coalgebra (X,a) is the usual one, with formulas of the form $\Diamond_q\phi$ being interpreted through

$$x \models \Diamond_q\phi \iff a(x)([\![\phi]\!]) \geqslant q.$$

From the results in Section 5.4, we get that this logic is expressive for $\mathfrak{S}_1$-coalgebras. On the other hand, it is easy to see that for every coalgebra (X,a) and every formula ϕ, we have

$$[\![\phi]\!]_{(X,a)} = X.$$

Thus, any two states of $\mathfrak{S}_1$ are logically equivalent, and hence behaviourally equivalent. This is not surprising, since the category of $\mathfrak{S}_1$-coalgebras has a terminal object, whose underlying space has precisely one element. This existence of such a terminal coalgebra can be established by direct calculation. It can also be derived from the observation that $\mathfrak{S}_1 1 \cong 1$ holds for any one-element measurable space 1. Hence, also, the limit of the terminal sequence of $\mathfrak{S}_1$ is 1.

Hence, from the view of coalgebraic logic, $\mathfrak{S}_1$-coalgebras are pretty uninteresting. But observe that we have a natural isomorphism

$$\mathfrak{S} \cong \mathfrak{S}_1(1 + (-)),$$

so $\mathfrak{S}_1$ can be used in place of $\mathfrak{S}$ in order to generate the measure-polynomial functors.

Moreover, as soon as we consider *models* over a non-empty set of variables, the situation changes and becomes more interesting. Of course, the logic is still expressive.

5.6.2. *Labelled Markov transition systems.* Let Act denote some non-empty set of actions. We are going to consider the functor $\mathfrak{S}^{\mathsf{Act}}$, whose coalgebras are *labelled Markov transition systems*, that is, measurable spaces equipped with an Act-indexed family of stochastic relations. In Section 4.2 we defined the family U^{Act} of $\mathfrak{S}^{\mathsf{Act}}$-predicate liftings as

$$U^{\mathsf{Act}} = \{\, \beta_{r,a} \mid r \in \mathbb{Q} \cap [0,1], a \in \mathsf{Act} \,\}, \qquad \beta_{r,a,X}(B) = \{\, (\mu_i)_{i\in\mathsf{Act}} \mid \mu_a(B) \geqslant r \,\}.$$

Writing $\pi_a : \mathfrak{S}^{\mathsf{Act}} \to \mathfrak{S}$ for the ath projection ($a \in \mathsf{Act}$), we have

$$\beta_{r,a,X}(B) = (\pi_{a,X})^{-1}[\beta_{r,X}(B)] = \mathfrak{B}\pi_{a,X} \cdot \beta_{r,X}(B),$$

hence U^{Act} is the lift of U along $(\pi_a)_{a\in\mathsf{Act}}$. From Proposition 5.15, we have that U^{Act} is expressive. Hence, the following logic is an expressive logic for Markov transition systems:

$$\phi ::= \top \mid \phi_1 \wedge \phi_2 \mid \langle a\rangle_q\phi$$

where, for a $\mathfrak{S}^{\mathsf{Act}}$-coalgebra (X,d), we have

$$x \models \langle a\rangle_q\phi \iff d(x)(a)([\![\phi]\!]) \geqslant q.$$

5.6.3. *Polyadic predicate liftings for Kripke models.* Let $\tau = (O, \mathsf{ar})$ be a modal similarity type. In Section 4.2, we associated with each modal operator $\Delta \in O$ of arity n and a rational number r, a predicate lifting $[\![\Delta_r]\!]$, which leads to the following interpretation of Δ_r in an $\mathfrak{S}(-)^n$-coalgebra (X,d):

$$(X,d), x \models \Delta_r(\phi_1,\ldots,\phi_n) \iff d(x)([\![\phi_1]\!] \times \cdots \times [\![\phi_n]\!]) \geqslant r.$$

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Clearly, these predicate liftings are generating and separating. Thus, we also have the product over all Δ_r, where $\Delta \in O$ and $r \in \mathbb{Q} \cap [0,1]$ is generating and separating for $\prod_{\Delta \in O} \mathfrak{S}(X^{\mathrm{ar}\Delta})$-coalgebras. The logic induced by this set of predicate liftings is just the extended modal logic introduced in Section 2.1, so this logic is expressive.

5.6.4. *Stochastic variants of CTL.* In model checking, one verifies properties of systems, often modelled as coalgebras, which are formulated using the notion of a computation path, that is, an infinite sequence of states in which each pair of consecutive states is related, in some suitable sense, by the coalgebra's dynamic.

Our notion of a computation path will incorporate residence times, so, formally, a computation path over a measurable space X is a sequence $(x_i, r_i)_{i \in \mathbb{N}}$ with $x_i \in X$, $r_i \in \mathbb{R}_{\geqslant 0} := \{\, r \in \mathbb{R} \mid r \geqslant 0 \,\}$. We equip the space $(X \times \mathbb{R}_{\geqslant 0})^{\mathbb{N}}$ with the product σ-algebra.

The logic CSL (*continuous stochastic logic*) is defined by the grammar

$$\phi ::= p \mid \top \mid \phi_1 \wedge \phi_2 \mid P_{\bowtie r}(N\phi) \mid P_{\bowtie r}(N^{[t,u]}\phi) \mid P_{\bowtie r}(\phi_1 U \phi_2) \mid P_{\bowtie r}(\phi_1 U^{[t,u]} \phi_2),$$

with $p \in \mathsf{Var}$, $\bowtie \in \{\, <, \leqslant, >, \geqslant \,\}$ and $r, t, u \in \mathbb{Q} \cap [0,1]$. We assume Var to be countable. Traditionally, CSL is presented using two different sets of formulas: *state-* and *path-formulas* (Clarke *et al.* 1999; Baier *et al.* 2003; Doberkat 2007b). However, the version of CSL we present here makes use of the steady-state operator and of negation, but adding these presents no difficulty. On the other hand, Doberkat (2007a), Doberkat (2007b) and Doberkat (2008) do not make use of the operators $P_{\bowtie r}(N-)$ and $P_{\bowtie r}(-U-)$.

The operators $P_{\bowtie r}(N-)$ and $P_{\bowtie r}(N^{[t,u]}-)$ are variants of the *next* operator. Informally, $x \models P_{\bowtie r}(N\phi)$ should hold if the probability to reach, in one step, a state in which ϕ holds is $\bowtie r$, where $\bowtie \in \{\, <, \leqslant, >, \geqslant \,\}$. Similarly, $x \models P_{\bowtie r}(N^{[t,u]}\phi)$ holds if the probability to reach a state in which ϕ holds with a residence time in $[t,u]$ is $\bowtie r$. $P_{\bowtie r}(-U-)$ and $P_{\bowtie r}(-U^{[t,u]}-)$ are stochastic variants of the until-operator. The informal semantics of these operators reads as follows: $P_{\bowtie r}(\phi U \psi)$ holds in a state s if the probability that ψ holds in a state s' on some path starting in s such that in all states on this path before s' the formula ϕ holds is $\bowtie r$; this operator does not take residence times into account. $P_{\bowtie r}(\phi U^{[t,u]} \psi)$ holds in a state s if the probability to reach a state in which ψ holds in the time interval $[t,u]$ such that in all intermediate states ϕ holds is $\bowtie r$.

We will give an interpretation of CSL using the functor $\mathfrak{M} : \mathbf{Meas} \to \mathbf{Meas}$ defined by

$$\mathfrak{M}(X) = \mathfrak{S}((X \times \mathbb{R}_{\geqslant 0})^{\mathbb{N}}),$$

with the obvious action on morphisms. We need some notation before we can give a coalgebraic interpretation for the operators of CSL. For a path

$$p = (x_i, r_i)_{i \in \mathbb{N}} \in (X \times \mathbb{R}_{\geqslant 0})^{\mathbb{N}}$$

and $t \in \mathbb{R}_{\geqslant 0}$, we write $p@t$ for x_i with i the least index for which $t < \sum_{j=0}^{i} t_j$ holds. Thus, $p@t$ is the state in which the path is at time t. The modal operators of CSL are

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

interpreted using the following predicate liftings for $\mathfrak{M}$:

$$\begin{aligned}
[\![P_{\bowtie r}(N-)]\!]_X(A) &:= \{\, \mu \in \mathfrak{M}X \mid \mu\big(\{\, ((x_i,r_i))_{i\in\mathbb{N}} \mid x_1 \in A \,\}\big) \bowtie r \,\} \\
[\![P_{\bowtie r}(N^{[t,u]}-)]\!]_X(A) &:= \{\, \mu \in \mathfrak{M}X \mid \mu\big(\{\, ((x_i,r_i))_{i\in\mathbb{N}} \mid x_1 \in A,\, r_0 \in [t,u] \,\}\big) \bowtie r \,\} \\
[\![P_{\bowtie r}(-U-)]\!]_X(A,B) &:= \\
&\quad \{\, \mu \in \mathfrak{M}X \mid \mu\big(\{\, ((x_i,r_i))_{i\in\mathbb{N}} \mid \exists k : x_k \in B \wedge \forall 0 \leqslant i < k : x_i \in A \,\}\big) \bowtie r \,\} \\
[\![P_{\bowtie r}(-U^{[t,u]}-)]\!]_X(A,B) &:= \\
&\quad \{\, \mu \in \mathfrak{M}X \mid \mu\big(\{\, p \mid \exists s \in [t,u] : p@s \in B \wedge \forall 0 \leqslant s' < s : p@s' \in A \,\}\big) \bowtie r \,\},
\end{aligned}$$

with $p \in (X \times \mathbb{R}_{\geqslant 0})^{\mathbb{N}}$ for the last formula. Well-definedness of these predicate liftings follows from the results in Doberkat (2007b, Section 6.3).

We write $\mathfrak{R} : \mathbf{Meas} \to \mathbf{Meas}$ for the functor given by $X \mapsto \mathfrak{S}_1(X \times \mathbb{R}_{\geqslant 0})$ with the obvious action on morphisms; recall that $\mathfrak{S}_1$ is the subfunctor of $\mathfrak{S}$ given by the probabilities, that is, those measures $\mu \in \mathfrak{S}X$ with $\mu(X) = 1$. Thus, an $\mathfrak{R}$-coalgebra is a stochastic relation that incorporates residence times. From the perspective of model checking, we are less interested in general $\mathfrak{M}$-coalgebras than in those coalgebras that arise as path-coalgebras of $\mathfrak{R}$-coalgebras in the sense of the following proposition.

Proposition 5.30. Let $(X, a : X \to \mathfrak{S}_1(X \times \mathbb{R}_{\geqslant 0}))$ be an $\mathfrak{R}$-coalgebra with X a Standard Borel space. Then there exists a $\mathfrak{M}$-coalgebra (X, b) with

$$a(x)(B) = b(x)\big(B \times \prod_{i\geqslant 1} X \times \mathbb{R}_{\geqslant 0}\big)$$

for all $x \in X$, $B \in \mathfrak{B}(X \times \mathbb{R}_{\geqslant 0})$. We call (X, b) the *path coalgebra* of (X, a).

Proof. See Doberkat (2009b, Proposition 2.4.3). □

Using the path coalgebra construction, it is possible to interpret CSL in an $\mathfrak{R}$-coalgebra (X, a) with X Standard Borel: x satisfies ϕ in (X, a) if and only if x satisfies ϕ in the path-coalgebra.

We define for $\bowtie \in \{<, \leqslant, >, \geqslant\}$, $r, t \in \mathbb{Q}$ an $\mathfrak{R}$-predicate lifting $\kappa_{\bowtie r,t}$ by

$$\kappa_{\bowtie r,t,X}(A) = \{\, \mu \in \mathfrak{S}(X \times \mathbb{R}_{\geqslant 0}) \mid \mu(A \times [0,t]) \bowtie r \,\}.$$

It is not hard to see that the set

$$K := \{\, \kappa_{\bowtie r,t} \mid \bowtie \in \{<, \leqslant, >, \geqslant\},\, r \in \mathbb{Q} \cap [0,1], t \in \mathbb{Q} \,\}$$

is a separating set of predicate liftings. Moreover, we have, for each $\mathfrak{R}$-coalgebra (X, a), X Standard Borel and (X, b) its path-coalgebra,

$$x \models \phi \text{ in } (X,a) \iff x \models \tilde{\phi} \text{ in } (X,b),$$

with $\phi \in \mathscr{L}(K)$ and $\tilde{\phi}$ the CSL formula that arises from ϕ by replacing every $\langle \kappa_{\bowtie r,t} \rangle$ by $P_{\bowtie r}(N^{[0,t]}-)$. By means of this purely syntactic transformation, we can treat $\mathscr{L}(K)$ as a fragment of CSL.

Using this semantics of CSL in $\mathfrak{R}$-coalgebras, we can prove the following proposition.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Proposition 5.31. The logic CSL is expressive for models based on $\mathfrak{R}$-coalgebras over Standard Borel spaces.

Proof. Let $\mathsf{M} = (X, a, V)$ and $\mathsf{N} = (Y, b, W)$ be models with $a : X \to \mathfrak{R}X$ and $b : Y \to \mathfrak{R}Y$ and X, Y Standard Borel. If states x and y of M and N, respectively, satisfy exactly the same CSL formulas, then they also satisfy exactly the same $\mathscr{L}(K)$ formulas. Since K is separating and countable and X and Y are Standard Borel, Corollary 5.12 shows that x and y are behaviourally equivalent. □

Using Theorem 5.10, one can also show that logically equivalent models based on $\mathfrak{R}$-coalgebras over Standard Borel spaces are behaviourally equivalent.

6. Bisimilarity

This section is devoted to a discussion of the relationship between behavioural equivalence and bisimilarity for stochastic relations. It turns out that this question cannot be answered without topological assumptions for the spaces involved. Specifically, we require the spaces to be Polish, or, a bit more generally, to be their Borel images, that is, analytic spaces. These spaces were introduced in Section 3.4, together with some properties that are helpful for our purposes.

Once we know that semi-pullbacks exist for the category of stochastic relations over these spaces, we can perform a construction that yields a span from a cospan of stochastic relations. We will first introduce the necessary notions from topology and measure theory, including a powerful selection argument due to Kuratowski and Ryll-Nardzewski, and we establish an easy to use criterion for the existence of a measurable selector in Section 6.1.

The notion of bisimilarity for stochastic relations is introduced in Section 6.2 and investigated in the context of semi-pullbacks, the existence of which is established in Section 6.3 as an application and illustration of selection. Congruences are introduced in Section 6.4 after we have done some leg work dealing with smooth equivalence relations; the relationship between congruences and bisimulations (which has been stressed in universal coalgebra as well) is then investigated in Section 6.4.2, again through a selection argument, demonstrating the power of this proof technique. Finally, we define the behavioural equivalence of stochastic relations and show that two relations are behaviourally equivalent if and only if there exist congruences such that the corresponding factors are isomorphic. Having done this, we relate behavioural equivalence to bisimilarity: two relations are bisimilar if and only if they are behaviourally equivalent if and only if they have isomorphic factors, provided the congruences are non-trivial.

6.1. *Selections*

Let X be a metric space, Z be a Polish space and Γ be a set-valued map from X to the non-empty and closed subsets of $\mathfrak{S}(Z)$. We will show that under a mild condition we can find a measurable selector M for Γ, that is, a stochastic relation M such that $M(x) \in \Gamma(x)$ holds for all $x \in X$. This method will turn out to be surprisingly versatile for

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

our purposes. It rests on the following fairly general selection result due to Kuratowski and Ryll-Nardzewski (Srivastava 1998, Theorem 5.2.1).

Theorem 6.1. Given a measurable space $(X, \mathscr{A})$ and a Polish space Z, we assume that Γ maps X to the non-empty closed subsets of Z such that

$$\exists\Gamma(G) := \{\, x \in X \mid \Gamma(x) \cap G \neq \varnothing \,\}$$

is a member of $\mathscr{A}$ for each open set $G \subseteq Z$. Then there exists an $\mathscr{A}$-$\mathscr{B}(Z)$-measurable map $f : X \to Z$ with $f(x) \in \Gamma(x)$ for all $x \in X$.

Actually, the Kuratowski and Ryll-Nardzewski Theorem is a bit more general, but we do not need this generality here. We observe that the function $\exists\Gamma$ is just the predicate transformer introduced by Γ and the natural transformation $[\![\Diamond]\!]$; *cf.* Section 4.1. We apply Theorem 6.1 in a setting where the space of subprobabilities of a Polish space serves as the range of the set-valued map.

Proposition 6.2. Let X_i, Y_i be Polish spaces and $K_i : X_i \rightsquigarrow Y_i$ be a continuous stochastic relation, $i = 1, 2$. Let $A \subseteq X_1 \times X_2$ and $B \subseteq Y_1 \times Y_2$ be closed with projections equal to the base spaces, and assume that for $(x_1, x_2) \in A$ the set

$$\Gamma(x_1, x_2) := \{\, \mu \in \mathfrak{S}(B) \mid \mathfrak{S}(\beta_i)(\mu) = K_i(x_i),\ i = 1, 2 \,\}$$

is not empty, with $\beta_i : B \to Y_i$ denoting the projections. Then there exists a stochastic relation $M : A \rightsquigarrow B$ such that $M(x_1, x_2) \in \Gamma(x_1, x_2)$ for all $(x_1, x_2) \in A$.

Proof.

(1) Let $\overline{Y_i}$ for $i = 1, 2$ be the closure of Y_i in the Hilbert cube and $\overline{B}$ be the closure of B in $\overline{Y}_1 \times \overline{Y}_2$. Then $\overline{B}$ is compact and contains the embedding of B into $\overline{Y}_1 \times \overline{Y}_2$, which we identify with B, as a Borel subset. This is the case since Y_i is a Borel subset in its compactification. The projections $\overline{\beta}_i : \overline{B} \to \overline{Y}_i$ are the continuous extensions to the projections $\beta_i : B \to Y_i$.

(2) The map $r_i : \mathfrak{S}(Y_i) \to \mathfrak{S}\left(\overline{Y}_i\right)$ with $r_i(\mu)(G) := \mu(G \cap Y_i)$ for $G \in \mathscr{B}(\overline{Y}_i)$ is continuous (in fact, it is an isometry with respect to the respective Hutchinson metrics). For $(x_1, x_2) \in A$, we define the set

$$\Gamma_0(x_1, x_2) := \{\, \mu \in \mathfrak{S}\left(\overline{B}\right) \mid \mathfrak{S}\left(\overline{\beta}_i\right)(\mu) = (r_i \cdot K_i)(x_i),\ i = 1, 2 \,\}.$$

Thus Γ_0 maps A to the non-empty closed subsets of $\mathfrak{S}\left(\overline{B}\right)$ since $\mathfrak{S}\left(\overline{\beta}_i\right)$ and $r_i \cdot K_i$ are continuous for $i = 1, 2$. If $\mu \in \Gamma_0(x_1, x_2)$, then

$$\begin{aligned}
\mu(\overline{B} \setminus B) &\leqslant \mu\big(\overline{B} \cap (\overline{Y}_1 \setminus Y_1 \times \overline{Y}_2) \cup (\overline{Y}_1 \times \overline{Y}_2 \setminus Y_2)\big) \\
&= \mathfrak{S}\left(\overline{\beta}_1\right)(\mu)(\overline{Y}_1 \setminus Y_1) + \mathfrak{S}\left(\overline{\beta}_2\right)(\mu)(\overline{Y}_2 \setminus Y_2) \\
&= \big(r_1 \cdot K_1\big)(x_1)(\overline{Y}_1 \setminus Y_1) + \big(r_2 \cdot K_2\big)(x_2)(\overline{Y}_2 \setminus Y_2) \\
&= 0.
\end{aligned}$$

Hence all members of $\Gamma_0(x_1, x_2)$ are concentrated on B.

(3) Let $C \subseteq \mathfrak{S}\left(\overline{B}\right)$ be compact, and assume that $(t_n)_{n \in \mathbb{N}}$ is a converging sequence in A with $t_n \in \exists\Gamma_0(C)$ for all $n \in \mathbb{N}$ such that $t_n \to t \in A$. Then there exists $\mu_n \in C \cap \Gamma_0(t_n)$.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Since C is compact, there exists a convergent subsequence, which we assume to be the sequence itself, so $\mu_n \rightharpoonup_w \mu$ for some $\mu \in C$. Continuity of $\mathfrak{S}\left(\overline{\beta}_i\right)$ and $K_i(x_i)$ for $i = 1, 2$ implies $\mu \in \Gamma_0(\mathrm{t})$. Consequently, $\exists\Gamma_0(C)$ is a closed subset of A.

(4) Since $\mathfrak{S}\left(\overline{B}\right)$ is compact by Theorem 3.15, we may represent each open set G as a countable union of compact sets $(C_n)_{n\in\mathbb{N}}$, such that

$$\exists\Gamma_0(G) = \bigcup_{n\in\mathbb{N}} \exists\Gamma_0(C_n),$$

hence $\exists\Gamma_0(G)$ is a Borel set in A. Theorem 6.1 gives us a stochastic relation $M_0 : A \rightsquigarrow \overline{B}$ with $M_0(x_1, x_2) \in \Gamma_0(x_1, x_2)$ for all $(x_1, x_2) \in A$. We then define $M(x_1, x_2)$ as the restriction of $M_0(x_1, x_2)$ to the Borel sets of B, so $M : A \rightsquigarrow B$ is the desired relation because $M_0(x_1, x_2)(\overline{B} \setminus B) = 0$. □

6.2. *Bisimilarity*

Bisimulations for stochastic relations are introduced as spans of morphisms such that common events exist. They relate two systems in terms of their elements, and hence in terms of non-deterministic relations of their state spaces. We will look at the general coalgebraic case first.

The discussion in Section 2.3 motivates calling two coalgebras bisimilar if and only if they are related through a span of morphisms. Specifically, the coalgebras (X, a) and (Y, b) are said to be *bisimilar* if and only if there exists a span of morphisms, that is, there exists a coalgebra (Z, c) and morphisms

$$(X, a) \xleftarrow{\ f\ } (Z, c) \xrightarrow{\ f'\ } (Y, b)\ .$$

The coalgebra (Z, c) in the middle is sometimes called *mediating*.

Bisimilarity for stochastic relations is defined in this spirit, bearing in mind that things are a bit more involved than for coalgebras. While we only have to take into account one state space for a coalgebra, a stochastic relation operates on two spaces. Nevertheless, we can use the definition for coalgebras as a guideline.

First we fix a suitable notion of morphism for stochastic relations. Let $\mathsf{K} = (X, Y, K)$ with $K : X \rightsquigarrow Y$ and $\mathsf{L} = (V, W, L)$ with $L : V \rightsquigarrow W$ be stochastic relations. Recall from Definition 3.10 that this means that $K : X \to \mathfrak{S}(Y)$ and $L : V \to \mathfrak{S}(W)$ are Borel maps. A *morphism* $\mathsf{K} \to \mathsf{L}$ is a pair $\mathsf{f} = (f, g)$ of Borel maps $f : X \to V$ and $g : Y \to W$ such that

$$L \cdot f = \mathfrak{S}(g) \cdot K$$

holds. This expands to

$$L\big(f(x)\big)(G) = K(x)\big(g^{-1}\big[G\big]\big)$$

for each $x \in X$ and each Borel set $G \in \mathfrak{B}W$.

Definition 6.3. The stochastic relations $\mathsf{K} = (X, Y, K)$ and $\mathsf{L} = (V, W, L)$ are said to be *bisimilar* if and only if there exist a stochastic relation $\mathsf{M} = (A, B, M)$ and surjective morphisms $\mathsf{f} = (f, g) : \mathsf{M} \to \mathsf{K}, \mathsf{h} = (h, k) : \mathsf{M} \to \mathsf{L}$ such that the σ-algebra $g^{-1}\big[\mathfrak{B}Y\big] \cap$

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

$k^{-1}[\mathfrak{B}W]$ is non-trivial, that is, does not just contain $\varnothing$ and B. The relation M is said to be *mediating*.

The first condition on bisimilarity is in accordance with the general definition of bisimilarity of coalgebras above; it requires that f and h form a span of morphisms

$$\mathsf{K} \xleftarrow{\mathsf{f}} \mathsf{M} \xrightarrow{\mathsf{h}} \mathsf{L}.$$

Hence, the following diagram is supposed to commute

$$\begin{array}{ccccc} X & \xleftarrow{f} & A & \xrightarrow{h} & V \\ {\scriptstyle K}\downarrow & & \downarrow{\scriptstyle M} & & \downarrow{\scriptstyle L} \\ \mathfrak{S}(Y) & \xleftarrow{\mathfrak{S}(g)} & \mathfrak{S}(B) & \xrightarrow{\mathfrak{S}(k)} & \mathfrak{S}(W) \end{array}$$

Thus, for each $a \in A, D \in \mathfrak{B}Y, E \in \mathfrak{B}W$, the equalities

$$K\big(f(a)\big)(D) = \big(\mathfrak{S}(g) \cdot M\big)(a)(D) = M(a)\big(g^{-1}[D]\big)$$

and

$$L\big(h(a)\big)(E) = \big(\mathfrak{S}(k) \cdot M\big)(a)(E) = M(a)\big(k^{-1}[E]\big)$$

should be satisfied. The second condition, however, is special since it states that we can find an event $C^* \in \mathfrak{B}B$ that is common to both K and L in the sense that

$$g^{-1}[D] = C^* = k^{-1}[E]$$

for some $D \in \mathfrak{B}Y$ and $E \in \mathfrak{B}W$ such that both $C^* \neq \varnothing$ and $C^* \neq B$ (note that for $C^* = \varnothing$ or $C^* = B$ we can always take the empty and full sets, respectively). Given such a C^* with D and E from above, we get for each $a \in A$,

$$\begin{aligned} K(f(a))(D) &= M(a)(g^{-1}[D]) \\ &= M(a)(C^*) \\ &= M(a)(k^{-1}[E]) \\ &= L(h(a))(E). \end{aligned}$$

Thus the event C^* ties K and L together. Loosely speaking, $g^{-1}[\mathfrak{B}Y] \cap k^{-1}[\mathfrak{B}W]$ can be described as the σ-algebra of common events, which is required to be non-trivial. Note that without the second condition, two relations K and L that are strictly probabilistic (that is, for which the entire space is always assigned probability 1) would always be bisimilar. To see this, put $A := X \times V$ and $B := Y \times W$, and for $(x, v) \in A$, set $M(x, v) := K(x) \otimes L(v)$ as the mediating relation, that is, define M pointwise to be the product measure of K and L. Then the projections will make the diagram commutative. It is also clear that this argument does not work for the subprobabilistic case.

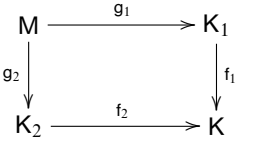
6.3. *Semi-pullbacks exist*

We want to prove that behaviourally equivalent relations are bisimilar – this will be established once we know that we can construct a span from a cospan of morphisms for

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

stochastic relations. This is what Proposition 6.4 states. A measurable space is said to be *separable* if and only if its σ-algebra is countably generated and separates points. It is known that $(Z, \mathscr{C})$ is a separable measurable space if and only if there exists a separable metric d on Z such that $\mathscr{C} = \mathscr{B}(Z, d)$.

Proposition 6.4. Let $\mathsf{K}_i = (X_i, Y_i, K_i)$ be stochastic relations over Polish spaces, and assume that $\mathsf{K} = (X, Y, K)$ is a stochastic relation, where X, Y are separable spaces. Assume that we have a cospan of surjective morphisms $\mathsf{f}_i : \mathsf{K}_i \to \mathsf{K}, i = 1, 2$. Then there exists a stochastic relation M and surjective morphisms $\mathsf{g}_i : \mathsf{M} \to \mathsf{K}_i, i = 1, 2$ rendering the following diagram commutative:



M is defined over Polish spaces.

Thus, one may say that semi-pullbacks exist in the category of stochastic relations between Polish spaces with surjective morphisms.

The proof of this statement is quite extensive and cannot be given here in its entirety. It requires a measure theoretic extension argument – basically, we will be able to make sure that we can find a suitable object on a sub-σ-algebra, and the extension will then make sure that this partial solution can be extended to the entire σ-algebra. This extension result is established through some classical results in mathematical analysis.

Proposition 6.5. Let X_1 and X_2 be Polish spaces and $(Z, \mathscr{C})$ be a separable measurable space, and assume that

$$g_i : X_i \to Z \qquad (i = 1, 2)$$

are measurable and surjective maps. Define

$$S := \{ (x_1, x_2) \in X_1 \times X_2 \mid g_1(x_1) = g_2(x_2) \},$$

endow S with the trace $\mathscr{B}(S)$ of the product σ-algebra, and assume that subprobability measures $\mu_1 \in \mathfrak{S}(X_1), \mu_2 \in \mathfrak{S}(X_2), \theta \in \mathfrak{S}(S)$ are given such that

$$\forall E_i \in g_i^{-1}[\mathscr{C}] : \mathfrak{S}(\pi_i)(\theta)(E_i) = \mu_i(E_i) \qquad (i = 1, 2)$$

holds, where $\pi_1 : S \to X_1, \pi_2 : S \to X_2$ are the projections. Then there exists $\mu \in \mathfrak{S}(S)$ such that

$$\forall E_i \in \mathscr{B}(X_i) : \mathfrak{S}(\pi_i)(\mu)(E_i) = \mu_i(E_i) \qquad (i = 1, 2)$$

holds.

Technically, this result will enable us to apply Proposition 6.2 to give a proof of Proposition 6.4.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Proof of Proposition 6.4.

(1) Assume $\mathsf{K}_i = (X_i, Y_i, K_i)$ with $\mathsf{f}_i = (f_i, g_i)$, $i = 1, 2$. Because of Proposition 3.12, we may assume that the respective σ-algebras on X_1 and X_2 are obtained from Polish topologies that render f_1 and K_1, as well as f_2 and K_2, continuous. These topologies are fixed for the proof. Put

$$A := \{ (x_1, x_2) \in X_1 \times X_2 \mid f_1(x_1) = f_2(x_2) \}$$
$$B := \{ (y_1, y_2) \in Y_1 \times Y_2 \mid g_1(y_1) = g_2(y_2) \}.$$

Both A and B are closed, and hence Polish. Let $\alpha_i : A \to X_i$ and $\beta_i : B \to Y_i$ denote the projections, $i = 1, 2$. The diagrams

$$\begin{array}{ccccc} X_1 & & \xrightarrow{f_1} & & X \xleftarrow{f_2} X_2 \\ {\scriptstyle K_1}\downarrow & & & & \downarrow{\scriptstyle K} \quad \downarrow{\scriptstyle K_2} \\ \mathfrak{S}(Y_1) & & \xrightarrow{\mathfrak{S}(g_1)} & & \mathfrak{S}(Y) \xleftarrow{\mathfrak{S}(g_2)} \mathfrak{S}(Y_2) \end{array}$$

are commutative by assumption, so we know for $x_i \in X_i$ that

$$K(f_i(x_i)) = \mathfrak{S}(g_i)(K_i(x_i)) \qquad (i = 1, 2)$$

holds. The construction implies that

$$(g_1 \cdot \beta_1)(y_1, y_2) = (g_2 \cdot \beta_2)(y_1, y_2)$$

is true for $(y_1, y_2) \in B$, and that $g_1 \cdot \beta_1 : B \to Y$ is surjective.

(2) Fix $(x_1, x_2) \in A$. We now use separability of the target spaces. We know that the image of a surjective map under $\mathfrak{S}$ is onto again (Doberkat 2007b), so there exists $\mu_0 \in \mathfrak{S}(B)$ with

$$\mathfrak{S}(g_1 \cdot \beta_1)(\mu_0) = K(f_1(x_1)).$$

Consequently,

$$\mathfrak{S}(g_i \cdot \beta_i)(\mu_0) = \mathfrak{S}(g_i)(K_i(x_i)), i = 1, 2).$$

But this means

$$\forall E_i \in g_i^{-1}[\mathscr{B}(Y)] : \mathfrak{S}(\beta_i)(\mu_0)(E_i) = K_i(x_i)(E_i) \qquad (i = 1, 2).$$

Put

$$\Gamma(x_1, x_2) := \{ \mu \in \mathfrak{S}(B) \mid \mathfrak{S}(\beta_1)(\mu) = K_1(x_1) \wedge \mathfrak{S}(\beta_2)(\mu) = K_2(x_2) \}.$$

Then Proposition 6.5 shows that $\Gamma(x_1, x_2) \neq \varnothing$.

(3) The set

$$\exists\Gamma(C) = \{ (x_1, x_2) \in A \mid \Gamma(x_1, x_2) \cap C \neq \varnothing \}$$

is closed in A for compact $C \subseteq \mathfrak{S}(B)$. This can be shown in exactly the same way as in the second part of the proof for Proposition 6.2 from which we now infer that there exists a measurable map $M : A \to \mathfrak{S}(B)$ such that $M(x_1, x_2) \in \Gamma(x_1, x_2)$ holds for every $(x_1, x_2) \in A$. Thus $M : A \rightsquigarrow B$ is a stochastic relation with

$$K_i \cdot \alpha_i = \mathfrak{S}(\beta_i) \cdot M \qquad (i = 1, 2).$$

Hence $\mathsf{M} := (A, B, M)$ is the desired semi-pullback. □

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

This construction differs considerably from that in Edalat (1999). Edalat constructs a selector for Γ explicitly through a disintegration argument. This selector, however, is not Borel measurable but rather universally measurable; this makes the results given in Desharnais *et al.* (2002) weaker than the ones obtained here. Recall that a set $C \subseteq Z$ in an analytic space Z is said to be *universally measurable* if and only if for each $\mu \in \mathfrak{S}(Z)$ there exist Borel sets C_1 and C_2 with $C_1 \subseteq C \subseteq C_2$ such that $\mu(C_2 \setminus C_1) = 0$. The universally measurable sets form a σ-algebra, which is called, you guessed it, the σ-algebra of universally measurable sets.

Edalat's construction may also be replaced by a selection argument: let $\mathscr{U}(A)$ be the σ-algebra of universally measurable sets constructed from the Borel sets of A, then

$$\Gamma \in \mathscr{U}(A) \otimes \mathscr{B}(\mathfrak{S}(B))$$

follows from

$$\Gamma \in \mathscr{B}(A \times \mathfrak{S}(B)) = \mathscr{B}(A) \otimes \mathscr{B}(\mathfrak{S}(B)).$$

The von Neumann Selection Theorem (Srivastava 1998, Theorem 5.5.7) then gives a universally measurable selector W for Γ, that is, a selector that is measurable with respect to the universally measurable sets. A universally measurable map has for any $\mu \in \mathfrak{S}(A)$ a Borel measurable modification that coincides with it except on a set of μ-measure 0 by Dedecker *et al.* (2004, Lemma 4.1), so we get from Edalat's construction that there exists a stochastic relation $W_\mu : A \rightsquigarrow B$ that is a selector μ-almost everywhere.

6.4. *Congruences and bisimulations*

With coalgebras there is a close relationship between congruences and bisimulations (Rutten 2000). This is also the case for stochastic relations, as we will now give ample opportunity to observe.

6.4.1. *Smooth relations.* Before defining congruences, we will first have a look at the underlying equivalence relations. We say an equivalence relation α on an analytic space X is *smooth* (or *countably generated*) if and only if there exists a sequence $(A_n)_{n\in\mathbb{N}}$ of Borel sets such that $x \,\alpha\, x'$ if and only if $x \in A_n \iff x' \in A_n$ holds for each $n \in \mathbb{N}$.

These equivalence relations are of interest in the present context because they arise from the interpretation of modal or coalgebraic logics with a countable number of formulas: ρ_{M} introduced in (7) is smooth provided there are countably many formulas, and extensions of formulas are Borel sets, which they usually are.

The following is a characterisation of smooth equivalence relations that we will use over and over again; for a proof and further discussion, see Doberkat (2009b, Chapter 1.7).

Proposition 6.6. Let X be an analytic space, and assume that α is an equivalence relation on X. Then the following conditions are equivalent:

(a) α is smooth.

(b) $\alpha = \ker(f) := \{(x, x') \mid f(x) = f(x')\}$ for some Borel map $f : X \to P$ with P a separable metric space.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

(c) X/α is an analytic space with $\mathscr{B}(X/\alpha) = \{D \mid \eta_\alpha^{-1}[D] \in \mathbf{\Sigma}(\mathscr{B}(X), \alpha)\}$.

Thus the factor space of a smooth equivalence relation is again analytic, and we have a fairly good handle on the Borel sets. Note that each Borel set $Q \in \mathscr{B}(X/\alpha)$ can be written as $Q = \eta_\alpha[R]$ for some α-invariant Borel set $R \in \mathbf{\Sigma}(\mathscr{B}(X), \alpha)$ – take $R := \eta_\alpha^{-1}[Q]$ and use the fact that R is α-invariant. The following observation is a simple demonstration, which will be useful later.

Lemma 6.7. Let X and Y be analytic spaces and assume that $f : X \to Y$ is a surjective Borel map. If $R \in \mathbf{\Sigma}(\mathscr{B}(X), \ker(f))$, then $f[R] \in \mathscr{B}(Y)$.

Proof. Because R is a Borel set, $f[R]$ is an analytic subset of Y. Since f is onto and R is $\ker(f)$-invariant, we can infer $f[X \setminus R] = Y \setminus f[R]$. Thus $f[R]$ is coanalytic, so $f[R]$ is a Borel set by Souslin's Theorem (Theorem 3.14). □

The following is a first simple observation on Borel isomorphisms.

Lemma 6.8. Let X and Y be analytic spaces and assume that $f : X \to Y$ is a surjective Borel map. Then $X/\ker(f)$ is isomorphic to Y.

Proof.

(1) Define $\widetilde{f}([x]_{\ker(f)}) := f(x)$. Then $\widetilde{f} : X/\ker(f) \to Y$ defines a bijection. Because $X/\ker(f)$ has the final σ-algebra with respect to $\eta_{\ker(f)}$, $\widetilde{f}$ is Borel since f is.
(2) Let Q be a Borel set in $X/\ker(f)$. We will show that $\widetilde{f}[Q]$ is a Borel set in Y. In fact, we know from Proposition 6.6 that $Q = \eta_{\ker(f)}[R]$ for some $R \in \mathbf{\Sigma}(\mathscr{B}(X), \ker(f))$, so it suffices to show that $f[R] \in \mathscr{B}(Y)$. But this is just what Lemma 6.7 says. □

6.4.2. *Congruences.* Turning to congruences proper, a *congruence* c for a stochastic relation $\mathsf{K} = (X, Y, K)$ for the analytic spaces X and Y is a pair (α, β) of smooth equivalence relations on X and Y, respectively, such that there exists a (unique) stochastic relation $\mathsf{K}/\mathsf{c} = (X/\alpha, Y/\beta, K_\mathsf{c})$ with

$$K_\mathsf{c} \cdot \eta_\alpha = \mathfrak{S}(\eta_\beta) \cdot K, \tag{*}$$

such that the following diagram commutes:

$$\begin{array}{ccc} X & \xrightarrow{\eta_\alpha} & X/\alpha \\ {\scriptstyle K}\downarrow & & \downarrow{\scriptstyle K_\mathsf{c}} \\ \mathfrak{S}(Y) & \xrightarrow{\mathfrak{S}(\eta_\beta)} & \mathfrak{S}(Y/\beta) \end{array}$$

An equivalent formulation states that

$$x \,\alpha\, x' \implies K(x)(B) = K(x')(B)$$

for all β-invariant Borel sets $B \subseteq Y$. Thus, if α cannot distinguish x and x', then $K(x)$ and $K(x')$ will assign identical probabilities to those sets that cannot be distinguished by

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

β. In fact, expanding the definition of a congruence, we see that equation (*) expands to

$$K_{\mathsf{c}}([x]_\alpha)(G) = K(x)(\eta_\beta^{-1}[G])$$

for each $x \in X$ and each Borel set $G \subseteq Y/\beta$. Now let $B \in \mathbf{\Sigma}(\mathscr{B}(Y), \beta)$ and assume that $x \,\alpha\, x'$. This implies

$$K(x)(B) = K_{\mathsf{c}}([x]_\alpha)(\eta_\beta[B]) = K_{\mathsf{c}}([x']_\alpha)(\eta_\beta[B]) = K(x')(B)$$

(note that $B = \eta_\beta^{-1}[\eta_\beta[B]]$ since B is β-invariant, so $\eta_\beta[B] \in \mathscr{B}(Y/\beta)$ – however, the image of a Borel set under a Borel map is usually not a Borel set). This discussion also shows that the converse also holds: if the implication above holds, then the diagram commutes, where K_{c} is defined in the obvious way. Thus, we have a fairly appealing and intuitively satisfying characterisation of congruences.

The interplay between morphisms and congruences is illustrated by the following result.

Proposition 6.9. Let $\mathsf{f} : \mathsf{K} \to \mathsf{L}$ be a morphism with $\mathsf{f} = (f, g)$ and f and g be surjective. Assume that K and L are based on analytic spaces. Then

$$\ker(\mathsf{f}) := (\ker(f), \ker(g))$$

is a congruence for K. Conversely, if $\mathsf{c} = (\alpha, \beta)$ is a congruence, then

$$\eta_{\mathsf{c}} := (\eta_\alpha, \eta_\beta)$$

is a morphism $\mathsf{K} \to \mathsf{K}/\mathsf{c}$.

Proof. Since the second part is immediate from the definition of a congruence, we will just consider the first part.

(1) Assume $x \,\alpha\, x'$ and H are β-invariant. Then

$$\begin{aligned} K(x)(H) &= K(x)(g^{-1}[g[H]]) \\ &= L(f(x))(g[H]) \\ &= L(f(x'))(g[H]) \\ &= K(x')(H) \end{aligned}$$

since f is a morphism $\mathsf{K} \to \mathsf{L}$.

(2) For $x \in X$ and the Borel set $B \in \mathscr{B}(Y/\ker(g))$, we define

$$K'([x]_{\ker(f)})(B) := K(x)(\eta_{\ker(g)}^{-1}[B]) = \left(\mathfrak{S}\left(\eta_{\ker(g)}\right) \cdot K\right)(x)(B).$$

Then part 1 shows that K' is well defined. The map $t \mapsto K'(t)(B)$ is Borel measurable for each $B \in \mathscr{B}(Y/\ker(g))$ because $X/\ker(f)$ carries the final σ-algebra with respect to the factor map $\eta_{\ker(f)}$. Consequently, $K' : X/\ker(f) \rightsquigarrow Y/\ker(g)$ is a stochastic relation. It is immediate from its definition that we have

$$K' \cdot \eta_{\ker(f)} = \mathfrak{S}\left(\eta_{\ker(g)}\right) \cdot K.$$

Thus $\ker(\mathsf{f})$ is a congruence. □

We can take this a bit further since each surjective morphism gives rise to an isomorphism, as one would expect.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Proposition 6.10. Let $\mathsf{f} : \mathsf{K} \to \mathsf{L}$ be a surjective morphism with K and L based on analytic spaces. Then $\mathsf{K}/\ker(\mathsf{f})$ is isomorphic to L.

Proof.

(1) Assume $\mathsf{K} = (X, Y, K)$, $\mathsf{L} = (U, V, L)$ and $\mathsf{f} = (f, g)$. Because f and g are onto, we infer from Lemma 6.8 that U is Borel isomorphic to $X/\ker(f)$ through $\widetilde{f}$, and V is Borel isomorphic to $Y/\ker(g)$ through $\widetilde{g}$. We have to establish that $(\widetilde{f}, \widetilde{g}) : \mathsf{K}/\ker(\mathsf{f}) \to \mathsf{L}$ is a morphism.

(2) Consider the following diagram:

$$\begin{array}{ccccc} X & \xrightarrow{\eta_{\ker(f)}} & X/\ker(f) & \xrightarrow{\widetilde{f}} & U \\ {\scriptstyle K}\downarrow & & \downarrow{\scriptstyle K_{\ker(\mathsf{f})}} & & \downarrow{\scriptstyle L} \\ \mathfrak{S}(Y) & \xrightarrow{\mathfrak{S}(\eta_{\ker(g)})} & \mathfrak{S}\left(Y/\ker(g)\right) & \xrightarrow{\mathfrak{S}(\widetilde{g})} & \mathfrak{S}(V) \end{array}$$

The outer diagram commutes since $(\widetilde{f} \cdot \eta_{\ker(f)}, \widetilde{g} \cdot \eta_{\ker(g)}) = (f, g)$ is a morphism. The leftmost square commutes by the definition of a congruence. Since $\eta_{\ker(f)}$ is onto, the rightmost square also commutes. This establishes the assertion on $(\widetilde{f}, \widetilde{g})$. □

Congruences will now help in establishing the relationship between behavioural equivalence and bisimilarity.

We first show that there is a very close coalgebraic connection between congruences and bisimulations. This is also an illustration of the 'Existence By Selection' proof principle, which turns out to be surprisingly versatile and powerful.

Proposition 6.11. Let $\mathsf{K} = (X, Y, K)$ be a stochastic relation, and assume that X and Y are analytic spaces. Assume that α and β are smooth equivalence relations on X and Y, respectively. Consider:

(i) (α, β) is a congruence for K.

(ii) There exists a stochastic relation $\mathsf{M} = (\alpha, \beta, M)$ making the following diagram commutative:

$$\begin{array}{ccccc} X & \xleftarrow{\pi_{1,X}} & \alpha & \xrightarrow{\pi_{2,X}} & X \\ {\scriptstyle K}\downarrow & & {\scriptstyle M}\downarrow & & \downarrow{\scriptstyle K} \\ \mathfrak{S}(Y) & \xleftarrow{\mathfrak{S}(\pi_{1,Y})} & \mathfrak{S}(\beta) & \xrightarrow{\mathfrak{S}(\pi_{2,Y})} & \mathfrak{S}(Y) \end{array}$$

Then (ii) $\Longrightarrow$ (i) always holds for analytic spaces. Moreover, if X and Y are Polish, the conditions are equivalent.

Proof.

(1) Assume that (ii) holds, and let $x \,\alpha\, x'$. If $B \in \mathbf{\Sigma}\big(\mathscr{B}(Y), \beta\big)$, it is clear that

$$\pi_{1,Y}^{-1}\left[B\right] \cap \beta = \pi_{2,Y}^{-1}\left[B\right] \cap \beta,$$

so

$$K(x)(B) = M(x, x')(\pi_{1,Y}^{-1}\left[B\right]) = M(x, x')(\pi_{2,Y}^{-1}\left[B\right]) = K(x')(B).$$

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

This establishes the first part.

(2) To establish (i) $\Longrightarrow$ (ii), assume that X and Y are Polish. Since α and β are smooth, by Proposition 6.6, we find separable metric spaces P and Q and Borel maps $f : X \to P$ and $g : Y \to Q$ such that $\alpha = \ker(f)$ and $\beta = \ker(g)$. By Proposition 3.12, we may, and do, assume that both f and g are continuous, and that $K : X \to \mathfrak{S}(Y)$ is continuous. Thus α and β are closed subsets of their respective Cartesian products. As in the proof of Proposition 6.4, we define the set

$$\Gamma(x, x') := \{\, \mu \in \mathfrak{S}(\beta) \mid \mathfrak{S}\left(\pi_{1,Y}\right)(\mu) = K(x),\ \mathfrak{S}\left(\pi_{2,Y}\right)(\mu) = K(x') \,\}$$

for $(x, x') \in \alpha$. Hence $\Gamma(x, x')$ is closed. However, before we can apply Proposition 6.2, we need to make sure that $\Gamma(x, x')$ is always non-empty. Granted that, we can use Proposition 6.2 to find a stochastic relation $M : \alpha \rightsquigarrow \beta$ with the desired properties.

(3) Fix $(x, x') \in \alpha$, and put $\nu_1 := K(x), \nu_2 := K(x')$. We define the measure

$$\theta_1(B \times B \cap \beta) := \nu_1(B)$$

on the σ-algebra

$$\mathscr{G} := \{\, B \times B \cap \beta \mid B \in \mathbf{\Sigma}\big(\mathscr{B}(Y), \beta\big) \,\} \subseteq \mathscr{B}(\beta)$$

(note that $\nu_1(B) = \nu_2(B)$ holds).

The fact that $\mathscr{G}$ is a σ-algebra was shown in Doberkat (2007b, Lemma 5.9). We now extend θ_1 to a measure μ on all of $\mathscr{B}(\beta)$, and infer from Proposition 6.5 that

$$\mathfrak{S}\left(\pi_{i,Y}\right)(\mu)(E_i) = \nu_i(E_i)\ (i = 1, 2)$$

holds for $E_1, E_2 \in \mathscr{B}(\beta)$. Thus $\Gamma(x, x') \neq \varnothing$. □

Consequently, we can construct, at least in a Polish space, a bisimulation from a congruence, provided β is not the universal relation. This is because, under this condition, $\mathbf{\Sigma}\big(\mathscr{B}(Y), \beta\big) \neq \{\, \varnothing, Y \,\}$, and the σ-algebra $\mathscr{G}$ constructed in the proof of Proposition 6.11 is contained in the σ-algebra of common events, which is supposed to be non-trivial for a bisimulation.

Proposition 6.11 yields the following corollary as a simple consequence.

Corollary 6.12. Given a morphism $(f, g) : \mathsf{K} \to \mathsf{L}$ with K defined over Polish spaces, there exists a stochastic relation $M : \ker(f) \rightsquigarrow \ker(g)$ turning the corresponding projections into morphisms.

Proof. $\ker(f, g)$ is a congruence for K by Proposition 6.9. □

It is interesting to compare Corollary 6.12 with the coalgebraic statement that kernel of morphisms are bisimulation equivalences (Rutten 2000, Proposition 5.7) (bisimulation equivalences are Rutten's equivalent for congruences – see Rutten (2000, Proposition 5.8)). Two proofs are offered, and both require that the functor underlying the coalgebra preserves weak pullbacks.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

This provides support for the conjecture that the nearly ubiquitous requirement in coalgebraic arguments that the functors preserve weak pullbacks may be replaced in a stochastic context by selection arguments.

6.5. *Behavioural equivalence*

We define the behavioural equivalence of stochastic relations in the same way as for models by the existence of a pair of surjective morphisms; *cf.* Definition 4.6. Then we will show that behavioural equivalence and bisimilarity are the same for stochastic relations over analytic spaces, using congruences and factoring as a tool.

Definition 6.13. Two stochastic relations K_1 and K_2 are said to be *behaviourally equivalent* if and only if there exists a co-mediating relation K and surjective morphisms $\mathsf{K}_1 \xrightarrow{\mathsf{f}_1} \mathsf{K} \xleftarrow{\mathsf{f}_2} \mathsf{K}_2$.

Assume $\mathsf{f}_i = (f_i, g_i), \mathsf{K}_i = (X_i, Y_i, K_i)$ and $\mathsf{K} = (X, Y, K)$. The condition says that we can find an element $x_2 \in X_2$ for each given $x_1 \in X_1$ such that $f_1(x_1) = f_2(x_2)$ with the property that for each $B \in \mathfrak{B}Y$ with $B_i = g_i^{-1}[B], i = 1, 2$, we have, since f_1 and f_2 are morphisms,

$$\begin{aligned} K_1(x_1)(B_1) &= K_1(x_1)(g_1^{-1}[B]) \\ &= K(f_1(x_1))(B) \\ &= K(f_2(x_2))(B) \\ &= K_2(x_2)(B_2). \end{aligned}$$

Hence, the probability of hitting the set $B_1 = g_1^{-1}[B]$ through K_1 starting from x_1 equals the probability of hitting the set $B_2 = g_2^{-1}[B]$ from x_2 through K_2 if the observations for x_1 and x_2 coincide.

6.5.1. *An auxiliary isomorphism.* To investigate behavioural equivalence, we fix surjective morphisms $\mathsf{K}_1 \xrightarrow{\mathsf{f}_1} \mathsf{K} \xleftarrow{\mathsf{f}_2} \mathsf{K}_2$ and assume that K, K_1, and K_2 are based on analytic spaces. We will show that $\mathsf{K}_1/\ker(\mathsf{f}_1)$ and $\mathsf{K}_2/\ker(\mathsf{f}_2)$ are isomorphic. This is done through two lemmas. We define the relations

$$\begin{aligned} \mathfrak{X} &:= \{ ([x_1]_{\ker(f_1)}, [x_2]_{\ker(f_2)}) \mid x_1 \in X_1, x_2 \in X_2, f_1(x_1) = f_2(x_2) \} \\ \mathfrak{Y} &:= \{ ([y_1]_{\ker(g_1)}, [y_2]_{\ker(g_2)}) \mid y_1 \in Y_1, y_2 \in Y_2, g_1(y_1) = g_2(y_2) \}. \end{aligned}$$

Lemma 6.14. $\mathfrak{X}$ is the graph of a Borel isomorphism $\xi : X_1/\ker(f_1) \to X_2/\ker(f_2)$ with inverse ζ, and $\mathfrak{Y}$ is the graph of a Borel isomorphism $\psi : Y_1/\ker(g_1) \to Y_2/\ker(g_2)$ with inverse ω.

Proof. We will only establish the properties of ξ since the second assertion is established in exactly the same way.

It is clear that $\mathfrak{X}$ is the graph of a total map $\xi : X_1/\ker(f_1) \to X_2/\ker(f_2)$, and that ξ is injective. Since f_1 is onto, ξ is onto. To establish measurability, let $G \in \mathscr{B}(X_2/\ker(f_2))$.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Then

$$\eta^{-1}_{\ker(f_1)}\big[\xi^{-1}\big[G\big]\big] = f_1^{-1}\big[f_2\big[\eta^{-1}_{\ker(f_2)}\big[G\big]\big]\big].$$

Because $\eta^{-1}_{\ker(f_2)}\big[G\big]$ is $\ker(f_2)$-invariant, $f_2\big[\eta^{-1}_{\ker(f_2)}\big[G\big]\big]$ is a Borel set in X by Lemma 6.7, so

$$f_1^{-1}\big[f_2\big[\eta^{-1}_{\ker(f_2)}\big[G\big]\big]\big] \in \mathscr{B}(X_1).$$

But this set is actually $\ker(f_1)$-invariant, so

$$\eta^{-1}_{\ker(f_1)}\big[\xi^{-1}\big[G\big]\big] \in \mathscr{B}(X_1/\ker(f_1)),$$

which establishes Borel measurability for ξ. Interchanging the roles of the maps f_1 and f_2 in the above argument shows that ζ is a Borel map too. □

The following technical remarks are needed to establish that we have indeed just defined an isomorphism.

Lemma 6.15. $(\xi, \psi) : \mathsf{K}_1/\ker(\mathsf{f}_1) \to \mathsf{K}_2/\ker(\mathsf{f}_2)$ is an isomorphism.

Proof. We know from Lemma 6.14 that both

$$\xi : X_1/\ker(f_1) \to X_2/\ker(f_2)$$

and

$$\psi : Y_1/\ker(g_1) \to Y_2/\ker(g_2)$$

are Borel isomorphisms, so we have to establish that

$$\begin{array}{ccc} X_1/\ker(f_1) & \xrightarrow{\ \xi\ } & X_2/\ker(f_2) \\ \Big\downarrow{\scriptstyle K_{1,\ker(\mathsf{f}_1)}} & & \Big\downarrow{\scriptstyle K_{2,\ker(\mathsf{f}_2)}} \\ \mathfrak{S}\big(Y_1/\ker(g_1)\big) & \xrightarrow{\ \mathfrak{S}(\psi)\ } & \mathfrak{S}\big(Y_2/\ker(g_2)\big) \end{array}$$

commutes.

(1) Let $Q \in \mathscr{B}(Y_2/\ker(g_2))$ be given. Define $\widetilde{Q} := g_2\big[\eta^{-1}_{\ker(g_2)}\big[Q\big]\big]$. So $\widetilde{Q} \in \mathscr{B}(W)$ by Lemma 6.7, and $g_2^{-1}\big[\widetilde{Q}\big] = \eta^{-1}_{\ker(g_2)}\big[Q\big]$. Moreover,

$$\psi^{-1}\big[Q\big] = \eta_{\ker(g_1)}\big[\widetilde{Q}\big].$$

(2) Given $x_1 \in X_1$, let $x_2 \in X_2$ such that

$$\xi([x_1]_{\ker(f_1)}) = [x_2]_{\ker(f_2)}.$$

Hence $f_1(x_1) = f_2(x_2)$. We then have

$$\begin{aligned} K_{2,\ker(\mathsf{f}_2)}(\xi([x_1]_{\ker(f_1)}))(Q) &= K_{2,\ker(\mathsf{f}_2)}([x_2]_{\ker(f_2)})(Q) \\ &= K_2(x_2)(\eta^{-1}_{\ker(\mathsf{g}_2)}\big[Q\big]) \\ &= K_2(x_2)(g_2^{-1}\big[\widetilde{Q}\big]) \\ &= \big(\mathfrak{S}(g_2)\cdot K_2\big)(x_2)(\widetilde{Q}) \\ &= K(f_2(x_2))(\widetilde{Q}) \end{aligned}$$

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

because f_2 is a morphism. But we know that $f_1(x_1) = f_2(x_2)$, so

$$\begin{aligned} K_{2,\ker(\mathsf{f}_2)}(\xi([x_1]_{\ker(f_1)}))(Q) &= K(f_2(x_2))(\widetilde{Q}) \\ &= K(f_1(x_1)(\widetilde{Q}) \\ &= \big(\mathfrak{S}(g_1)\cdot K_1\big)(x_1)(\widetilde{Q}). \end{aligned}$$

Consequently,

$$\begin{aligned} K_{2,\ker(\mathsf{f}_2)}(\xi([x_1]_{\ker(f_1)}))(Q) &= K_1(x_1)(g_1^{-1}[\widetilde{Q}]) \\ &= K_{1,\ker(\mathsf{f}_1)}([x_1]_{\ker(f_1)})(\eta_{\ker(g_1)}[g_1^{-1}[\widetilde{Q}]]) \\ &= K_{1,\ker(\mathsf{f}_1)}([x_1]_{\ker(f_1)})(\psi^{-1}[Q]) \\ &= \big(\mathfrak{S}(\psi)\cdot K_{1,\ker(\mathsf{f}_1)}\big)([x_1]_{\ker(f_1)})(Q), \end{aligned}$$

which completes the proof. □

Summarising, we have now established the following result.

Proposition 6.16. If the stochastic relations K_1 and K_2 are behaviourally equivalent, there exist congruences c_1 on K_1 and c_2 on K_2 such that $\mathsf{K}_1/\mathsf{c}_1$ and $\mathsf{K}_2/\mathsf{c}_2$ are isomorphic.

Note that another way of deriving Proposition 6.16 is to observe that if K is co-mediating with $\mathsf{K}_1 \xrightarrow{\mathsf{f}_1} \mathsf{K} \xleftarrow{\mathsf{f}_2} \mathsf{K}_2$, we obtain from Proposition 6.10 the diagram

$$\begin{array}{ccccc} \mathsf{K}_1 & & & & \mathsf{K}_2 \\ \Big\downarrow{\scriptstyle \eta_{\ker(\mathsf{f}_1)}} & & & & \Big\downarrow{\scriptstyle \eta_{\ker(\mathsf{f}_2)}} \\ \mathsf{K}_1/\ker(\mathsf{f}_1) & \xrightarrow{\widetilde{\mathsf{f}}_1} & \mathsf{K} & \xleftarrow{\widetilde{\mathsf{f}}_2} & \mathsf{K}_2/\ker(\mathsf{f}_2) \end{array}$$

with isomorphisms $\widetilde{\mathsf{f}}_1$ and $\widetilde{\mathsf{f}}_2$. Proposition 6.16 is easily obtained from this observation. On the other hand, we wanted to demonstrate the technique of chasing measurable sets along images and inverse images, which is also helpful for other constructions.

6.5.2. *Isomorphic factors.* We can now give a characterisation of behaviourally equivalent relations. We say a congruence (α, β) is *non-trivial* if and only if β is neither the diagonal nor the universal relation.

Theorem 6.17. Let K_1 and K_2 be stochastic relations over analytic spaces. Consider the statements:

(a) K_1 and K_2 are behaviourally equivalent.
(b) $\mathsf{K}_1/\mathsf{c}_1$ is isomorphic to $\mathsf{K}_2/\mathsf{c}_2$ for suitable congruences c_1 on K_1 and c_2 on K_2.
(c) K_1 and K_2 are bisimilar.

Then

(1) (a) $\Longleftrightarrow$ (b).
(2) (b) $\Longrightarrow$ (c) provided the base spaces for the relations are Polish and the congruences are non-trivial.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

(3) If the base spaces for the relations are Polish and the mediating object is compact with continuous morphisms, then (c) $\Longrightarrow$ (a) holds.

Proof.

(1) The implication (a) $\Longrightarrow$ (b) is obtained from Proposition 6.16, and (b) $\Longrightarrow$ (a) is immediate.

(2) We now assume that $\mathsf{K}_i = (X_i, Y_i, K_i)$ is defined over Polish spaces with $\mathsf{c}_i = (\alpha_i, \beta_i)$ as non-trivial congruences. Let $(\xi, \zeta) : \mathsf{K}_2/\mathsf{c}_2 \to \mathsf{K}_1/\mathsf{c}_1$ be an isomorphism. Define

$$U := \{\, (x_1, x_2) \mid [x_1]_{\alpha_1} = \xi([x_2]_{\alpha_2}) \,\}$$
$$V := \{\, (y_1, y_2) \mid [y_1]_{\beta_1} = \zeta([y_2]_{\beta_2}) \,\}.$$

We infer from the proof of Proposition 6.4 that we can find a stochastic relation $M : U \rightsquigarrow V$ such that

$$\mathsf{K}_1 \xleftarrow{\pi_1} (U, V, M) \xrightarrow{\pi_2} \mathsf{K}_2$$

are surjective morphisms, where π_1 and π_2 are morphisms consisting of the corresponding projections. We have to show that the σ-algebra of common events is not trivial. Because β_1 is neither the diagonal nor the universal relation, we can find $B_1 \in \mathbf{\Sigma}\big(\mathscr{B}(Y_1), \beta_1\big)$ such that $\varnothing \neq B_1 \neq Y_1$. Define $B_2 := (\eta_{\beta_2} \circ \zeta)^{-1}\big[B_1\big]$. So $B_2 \in \mathbf{\Sigma}\big(\mathscr{B}(Y_2), \beta_2\big)$ is such that

$$B := \{\, (y_1, y_2) \in V \mid y_1 \in B_1 \,\} = \{\, (y_1, y_2) \in V \mid y_2 \in B_2 \,\},$$

hence $\varnothing \neq B$ is a common event. Taking $y_1 \notin B_1$ for some $y_1 \in Y_1$, we have $[y_1]_{\beta_1} = \zeta\big([y_2]_{\beta_2}\big)$, so $(y_1, y_2) \notin B$, but $(y_1, y_2) \in V$. Thus $B \neq V$. Thus K_1 and K_2 are bisimilar.

(3) (Sketch) Let M be the mediating object. Without loss of generality, we can assume that K_i is isomorphic to M/c_i for some congruence c_i with $i = 1, 2$. We obtain from the proof of Doberkat (2007b, Proposition 5.45) congruences d_i on M/c_i such that $\big(\mathsf{M}/\mathsf{c}_1\big)/\mathsf{d}_1$ and $\big(\mathsf{M}/\mathsf{c}_2\big)/\mathsf{d}_2$ are isomorphic. Consequently, we obtain

$$\begin{array}{ccc}
\mathsf{K}_1 & & \mathsf{K}_2 \\
\downarrow & & \downarrow \\
\mathsf{M}/\mathsf{c}_1 & & \mathsf{M}/\mathsf{c}_2 \\
\downarrow & & \downarrow \\
\big(\mathsf{M}/\mathsf{c}_1\big)/\mathsf{d}_1 & \xrightarrow{\ \cong\ } & \big(\mathsf{M}/\mathsf{c}_2\big)/\mathsf{d}_2
\end{array}$$

This establishes behavioural equivalence. □

Clearly, the crucial technical point in the proof of Theorem 6.17 is the existence of the congruences d_1 and d_2. The underlying equivalence relations are constructed as the respective transitive closure of a reflexive, symmetric and closed subset of the corresponding factor spaces; the compactness assumption makes sure that this closure remains manageable in terms of measurability. For a closer inspection of this construction, see the discussion in Doberkat (2007b, Chapter 5) or Doberkat (2006). It is not clear that

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

the implication in question can be reversed – the construction is akin to a pushout (albeit without the uniqueness property), but since analytic spaces are not closed under pushouts (Doberkat 2010a), there is reason to be suspicious.

Summarising, we have established that, under some not too restrictive conditions, the equivalence of behavioural equivalence and the existence of isomorphic factor spaces – note that no logic is needed as an intermediary enabling instance (hence logical equivalence does not enter the discussion at this point). We have also given an indication that bisimilarity implies behavioural equivalence, but this depends on a condition that appears to be somewhat restrictive, at least when the co-mediating object is required to be based on analytic spaces. When logic enters on the scene, however, this picture changes, as we saw in Example 5.11.

7. Further research

The current paper has presented an approach to the interpretation of modal logics and coalgebraic logic (their generalisation) through Kripke models on measurable spaces, thereby taking a broad variety of functors into account. This approach is fairly general since it makes a uniform line of argument available, which in turn abstracts away from specific details. These details may be of interest for particular logics, but they tend to obscure the broader picture. For example, when investigating behavioural equivalence for modal logics and for continuous time stochastic logics, particular idiosyncrasies of the respective functors had to be taken into account, but essentially the same arguments had to be applied. Using a functorial formulation and representing the modal operators through predicate liftings, that is, particular natural transformations, the mathematical essence of behavioural equivalence became much clearer. In this way, both the functor and the modal operators, now known as predicate liftings, turn into parameters for the logic. The functor, when considered in isolation, will determine the coalgebra upon which a Kripke model is based – the predicate liftings, which are also formulated in dependency on the functor, determine the face of the logic, its syntax and its interpretation.

This observation opens up a large area for future investigations. The topological case is touched on in Doberkat (2008), Doberkat and Schubert (2009) and Doberkat (2009b); here the basic scenario is provided by Polish spaces, and by analytic spaces as their continuous images. Polish spaces have pleasant measure theoretic properties, part of which was capitalised on in these papers, and they make the treatment of bisimulation in this context feasible through measurable selections. However, measurable spaces were missing, in general, so the challenge is to find a class of measurable spaces that is, on the one hand, sufficiently general, but, on the other hand, specific enough for us to derive interesting properties. It may well be that complete measurable spaces are an interesting class to look at: each measurable space can be completed, and these spaces are closed under the Souslin operation. Basing Kripke models on these spaces will certainly permit important insights into the structure of the stochastic interpretation of coalgebraic logics.

Jacobs and Sokolova (2010) showed that the extended basic logic is expressive for coalgebras of the subprobability functor over general measurable spaces. In order to

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

establish this result, they used Klin's approach to coalgebraic logic (Klin 2007). In this setting, a family of predicate liftings is replaced by a single natural transformation. The main technical requirement in Jacobs and Sokolova (2010) is that this single natural transformation in the category of measurable spaces is pointwise an initial monomorphism. It turns out that the induced natural transformation of a family of predicate liftings is pointwise injective if the set of predicate liftings is separating in the sense of the current paper. Moreover, the induced natural transformation is pointwise initial if the set of predicate liftings is generating in the sense of this paper. We do not know whether the converse implications hold. In any case, the technical requirements of Jacobs and Sokolova (2010) and the current paper are closely related, and this relationship warrants further investigation.

We should also mention some work on deduction systems for coalgebras over general measurable spaces. Goldblatt (2010) develops a theory of deduction for measurable polynomial functors, pointing out the proof theoretical side of the model theoretic approach developed by Moss and Viglizzo (Moss and Viglizzo 2006). Goldblatt obtains a proof-theoretic characterisation of the consequence relation for the logic of a measurable polynomial functor. Two points in his investigation deserve comment. First, his approach assumes the Lindenbaum property rather than proves it (he then proceeds to show that the deduction system is the least deduction system with this property, and the only one which is sound). Second, he describes σ-additivity axiomatically through his Countable Additivity Rule (in this way making sure that he does in fact construct a stochastic relation). Zhou, also starting from the work by Moss and Viglizzo, develops a sound and (weakly) complete probability logic for type spaces in a formal propositional language, and shows its decidability (Zhou 2009). In addition, he establishes negative results on strong completeness for two important types of sets of probability formulas. From these results, he went on to investigate a theory of probability logic that is parallel to that concerning belief types in game theory (Zhou 2010); in contrast to the work mentioned above, finite additivity suffices for this approach.

In the current paper, we have used a measure theoretic framework for modelling stochastic phenomena. Other approaches seem possible. For instance, one could use valuations on topological spaces or domains (Jones 1989). We think an investigation of these structures from the point of view of coalgebraic logic might be fruitful. Another possible line for further research comes from dropping the requirement that our measures are σ-additive. First steps in this direction were undertaken in Desharnais *et al.* (2009), where subprobability measures were replaced by super-additive functionals. It seems that the expressivity of simple modal languages can be retained, but one has to enlarge the syntactic part of the logic to incorporate countably infinite conjunctions and negation. This also warrants further investigation.

Acknowledgements

The authors are grateful for the comments and suggestions offered by the referee, which helped us to improve the presentation. We are also grateful to Ingo Battenfeld for his comments.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

References

Baier, C., Haverkort, B., Hermanns, H. and Katoen, J.-P. (2003) Model-checking algorithms for continuous time Markov chains. *IEEE Trans. Softw. Eng.* **29** (6) 524–541.

Billingsley, P. (1968) *Convergence of Probability Measures*, first edition, John Wiley and Sons.

Billingsley, P. (1995) *Probability and Measure*, John Wiley and Sons.

Blackburn, P., de Rijke, M. and Venema, Y. (2001) *Modal Logic*, Cambridge Tracts in Theoretical Computer Science **53**, Cambridge University Press.

Cîrstea, C. (2004) On logics for coalgebraic simulation. *Electronic Notes in Theoretical Computer Science* **106** 63–90.

Cîrstea, C. (2006) Modularity in coalgebra. *Electronic Notes in Theoretical Computer Science* **164** 3–26.

Cîrstea, C. and Pattinson, D. (2004) Modular construction of modal logics. In: Proceedings CONCUR 2004. *Springer-Verlag Lecture Notes in Computer Science* **2170** 258–275.

Clarke, E. M., Grumberg, O. and Peled, D. A. (1999) *Model Checking*, The MIT Press.

Danos, V., Desharnais, J., Laviolette, F. and Panangaden, P. (2006) Bisimulation and cocongruence for probabilistic systems. *Information and Computation* **204** 503–523.

Dedecker, J., Prieur, C. and de Fitte, P. (2004) Parametrized Kantorovich–Rubinštein theorem and application to the coupling of random variables. arXiv:math/0410052v2 [math PR].

Desharnais, J., Edalat, A. and Panangaden, P. (2002) Bisimulation of labelled Markov-processes. *Information and Computation* **179** (2) 163–193.

Desharnais, J., Laviolette, F. and Turgeon, A. (2009) A demonic approach to information in probabilistic systems. In: Bravetti, M. and Zavattaro, G. (eds.) CONCUR 2009 – Concurrency Theory. *Springer-Verlag Lecture Notes in Computer Science* **5710** 289–304.

Doberkat, E.-E. (2005) Semi-pullbacks for stochastic relations over analytic spaces. *Mathematical Structures in Computer Science* **15** 647–670.

Doberkat, E.-E. (2006) Stochastic relations: congruences, bisimulations and the Hennessy–Milner theorem. *SIAM J. Computing* **35** (3) 590–626.

Doberkat, E.-E. (2007a) The Hennessy–Milner equivalence for continuous-times stochastic logic with mu-operator. *J. Appl. Logic* **35** 519–544.

Doberkat, E.-E. (2007b) *Stochastic Relations. Foundations for Markov Transition Systems*, Chapman and Hall.

Doberkat, E.-E. (2008) Stochastic coalgebraic logic: Bisimilarity and behavioral equivalence. *Ann. Pure Appl. Logic* **155** 46–68.

Doberkat, E.-E. (2009a) Behavioral and logical equivalence of stochastic Kripke models in general measurable spaces. In: Chen, J. and Cooper, S. (eds.) TAMC 2009. *Springer-Verlag Lecture Notes in Computer Science* **5532** 192–200.

Doberkat, E.-E. (2009b) *Stochastic Coalgebraic Logic*, EATCS Monographs in Theoretical Computer Science, Springer-Verlag.

Doberkat, E.-E. (2010a) Lattice properties of congruences for stochastic relations. *Ann. Pure Appl. Logic* (in press).

Doberkat, E.-E. (2010b) Quotients for the Kleisli category over the Giry monad. Technical Report 182, Chair for Software Technology, Technische Universität Dortmund.

Doberkat, E.-E. and Schubert, C. (2009) Coalgebraic logic for stochastic right coalgebras. *Annals of Pure and Applied logic* **159** 268–284.

Dunford, N. and Schwartz, J. T. (1957) *Linear Operators*, Volume I, Interscience.

Edalat, A. (1999) Semi-pullbacks and bisimulation in categories of Markov processes. *Mathematical Structures in Computer Science* **9** (5) 523–543.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Edgar, G. A. (1998) *Integral, Probability and Fractal Measures*, Springer-Verlag.

Elstrodt, J. (1999) *Maß- und Integrationstheorie*, second edition, Springer-Verlag.

Goldblatt, R. (2010) Deduction systems for coalgebras over measurable spaces. *Journal of Logic and Computation* **20** (5) 1069–1100.

Halmos, P. R. (1950) *Measure Theory*, Van Nostrand Reinhold.

Jacobs, B. and Sokolova, A. (2010) Exemplaric expressivity of modal logics. *Journal of Logic and Computation* **20** (5) 1041–1068.

Jones, C. (1989) *Probabilistic Non-determinism*, Ph.D. thesis, University of Edinburgh.

Kechris, A. S. (1994) *Classical Descriptive Set Theory*, Graduate Texts in Mathematics, Springer-Verlag.

Klin, B. (2007) Coalgebraic modal logic beyond sets. *Electronic Notes in Theoretical Computer Science* **173** 177–201.

Kuratowski, K. (1966) *Topology*, Volume I, PWN – Polish Scientific Publishers and Academic Press.

Larsen, K. G. and Skou, A. (1991) Bisimulation through probabilistic testing. *Information and Computation* **94** 1–28.

Mac Lane, S. (1997) *Categories for the Working Mathematician*, Graduate Texts in Mathematics, Springer-Verlag.

Moss, L. and Viglizzo, I. (2006) Final coalgebras for functors on measurable spaces. *Information and Computation* **204** 610–636.

Moss, L. M. (1999) Coalgebraic logic. *Ann. Pure Appl. Logic* **96** 277–317.

Parthasarathy, K. R. (1967) *Probability Measures on Metric Spaces*, Academic Press.

Pattinson, D. (2004) Expressive logics for coalgebras via terminal sequence induction. *Notre Dame J. Formal Logic* **45** (1) 19–33.

Rudin, W. (1974) *Real and Complex Analysis*, second edition, Tata McGraw-Hill.

Rutten, J. J. M. M. (2000) Universal coalgebra: a theory of systems. *Theoretical Computer Science* **249** (1) 3–80. (Special issue on modern algebra and its applications.)

Schröder, L. (2008) Expressivity of coalgebraic modal logic: the limits and beyond. *Theoretical Computer Science* **390** 230–247.

Schröder, L. and Pattinson, D. (2007) Modular algorithms for heterogeneous modal logics. In: Arge, L., Cachin, C., Jurdzinski, T. and Tarlecki, A. (eds.) Proceedings ICALP. *Springer-Verlag Lecture Notes in Computer Science* **4596** 459–471.

Schubert, C. (2007) Coalgebraic logic over analytic spaces. Technical Report 170, Chair for Software Technology, Technische Universität Dortmund.

Schubert, C. (2009a) Coalgebraic logic over measurable spaces: behavioral and logical equivalence. In: Chen, Y., Doberkat, E.-E. and Jung, A. (eds.) Proceedings of the Fifth International Symposium on Domain Theory (ISDT 2009). *Electronic Notes in Theoretical Computer Science* **257** 71–85.

Schubert, C. (2009b) Coalgebraic modal logic over measurable spaces: behavioral versus logical equivalence. Technical Report 179, Chair for Software Technology, Technische Universität Dortmund.

Schubert, C. (2009c) Final coalgebras for measure-polynomial functors. In Chen, J. and Cooper, B. (eds.) TAMC 2009. *Springer-Verlag Lecture Notes in Computer Science* **5532** 325–334.

Srivastava, S. M. (1998) *A Course on Borel Sets*, Graduate Texts in Mathematics, Springer-Verlag.

Viglizzo, I. D. (2005a) *Coalgebras over measurable spaces*, Ph.D. thesis, Indiana University.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press

Viglizzo, I. D. (2005b) Final sequences and final coalgebras for measurable spaces. In: Fiadeiro, J., Harman, N., Roggenbach, M., and Rutten, J. (eds.) Algebra and Coalgebra in Computer Science – Proceedings First International Conference, CALCO 2005. *Springer-Verlag Lecture Notes in Computer Science* **3629** 395–407.

Yamasaki, Y. (1985) *Measures on infinite dimensional spaces*, World Scientific.

Zhou, C. (2009) A complete deductive system for probability logic. *J. Logic Comp.* **19** (6) 1427–1454.

Zhou, C. (2010) Probability logics for finitely additive beliefs. *J. Logic, Language and Information* **19** (3) 247–282.

https://doi.org/10.1017/S0960129510000526 Published online by Cambridge University Press