Comparing and Constructing (Co)Inductive Predicates

Ruben Turkenburg

Institute for Computing and Information Sciences

RADBOUD UNIVERSITY PRESS

Radboud Dissertation Series

Comparing and Constructing (Co)Inductive Predicates

Ruben Turkenburg

This research has been supported by the NWO grant OCENW.M20.053 and has been carried out under the auspices of the research school IPA (Institute for Programming research and Algorithmics) and Radboud University.

IPA dissertation series 2026-10

Comparing and Constructing (Co)Inductive Predicates
Ruben Turkenburg

Radboud Dissertations Series
ISSN: 2950-2772 (Online); 2950-2780 (Print)

Published by RADBOUD UNIVERSITY PRESS
Postbus 9100, 6500 HA Nijmegen, The Netherlands
www.radbouduniversitypress.nl

Design: Ruben Turkenburg
Cover: Proefschrift-aio.nl | Guus Gijben
Printing: DPN Rikken/Pumbo

ISBN: 9789465152448
DOI: 10.54195/9789465152448
Free download at: https://doi.org/10.54195/9789465152448

RADBOUD
UNIVERSITY
PRESS

Comparing and Constructing (Co)Inductive Predicates

Proefschrift ter verkrijging van de graad van doctor
aan de Radboud Universiteit Nijmegen
op gezag van de rector magnificus prof. dr. J.M. Sanders,
volgens besluit van het college voor promoties
in het openbaar te verdedigen op

donderdag 18 juni 2026

om 10:30 uur precies

door

Ruben Theodorus Cornelis Turkenburg

geboren op 2 april 1997
te York, Verenigd Koninkrijk

Promotoren

Dr. Jurriaan Rot

Prof. dr. Herman Geuvers

Manuscriptcommissie

Prof. dr. Bart Jacobs (voorzitter)

Prof. dr. Helle Hvid Hansen (Rijksuniversiteit Groningen)

Prof. dr. Alexander Kurz (Chapman University, Verenigde Staten)

Prof. dr. Lutz Schröder (FAU Erlangen-Nürnberg, Duitsland)

Prof. dr. Ana Sokolova (Paris Lodron Universität Salzburg, Oostenrijk)

Acknowledgements

There's only really one way to start this, and that's by thanking my supervisor Jurriaan: thank you! It goes without saying (but I'll say it anyway), that the last four years would not have been possible without you and your support. When you approached me about doing a PhD back in 2021, I was excited but also unsure. I had little idea then how many great new experiences it would give me, how much it would allow me to travel, and how many new things I would learn and research with you and other great collaborators. Looking back, saying yes was the obvious decision.

Thanks also to Herman for your work as my (other) promotor. We perhaps didn't meet as often as we planned, but the meetings we did have and your advice were always very valuable.

Many thanks to my further collaborators, who helped make the work of this thesis possible. Thanks Harsh for the great visits to Sheffield, and the noodles and many cups of coffee. Thanks Clemens and Ezra for the nice times in Glasgow. Thanks, Franck for your many insightful comments. Thanks, Thorsten and Bálint for getting me involved with the work on trees. Thanks to all of you for the help and interesting discussions which contributed to this thesis and beyond. I hope we will have the chance to meet and work together again in the future.

Thanks to Shin-ya and Ichiro for hosting me during my time in Japan. It was a wonderful experience, and although our discussions have not (yet) turned into full collaborations, my visits in Kyoto and Tokyo were a great inspiration. Thanks also to Kazuki for the advice on where to visit and showing me some of the sights and great food in Tokyo.

Thanks to everyone in Nijmegen and at the department who made it such a lovely place to be and work, and thanks especially to those who made office 1.08 what it was. Thanks Thorsten, Kazuki, Pedro, and of course Loes and Marck. I'm going to miss all the discussions (serious or otherwise), the many trips for coffee and tea, and, perhaps most of all, the puzzles.

Thanks to Jana for the discussions about academic life and beyond, and for the fun (at least as fun as it can be) times making and grading exams.

Thanks to Alexander, Ana, Bart, Helle, and Lutz for reading this thesis, and for the valuable feedback. Thanks also to you, and Henning and Yde for being on my defence committee.

Thanks to Harry, George, Sam A, Sam B, Richard, and everyone else I got to see on my trips back to the UK. I hope I can be around for many more York traditions in future. Thanks to Cass, Herman and Wessel for the walks in the park and accompanying discussions, and the dinners together in Nijmegen. I hope we get the chance to meet up again soon.

And finally to my family, Mum, Dad, Laura, Chris, Anna, Isaac, thank you for everything. Thanks for the many discussions, the help with proofreading this thesis and giving me places which will always feel like home.

Contents

CHAPTER 1

Introduction

Words make you think.
Music makes you feel.
A song makes you feel a thought.

E. Y. Harburg

In the study of computation, state-based systems form useful models of the possible behaviours of computers, the programs that run on them, and protocols defining how they communicate (among others). Such models are abstract mathematical objects, which allow us to focus on the most important properties of systems of interest. For instance, when a human interacts with a computer, we may be interested just in the inputs from the user and the outputs provided by the computer, and not precisely how the computer processes the inputs and produces the outputs. When trying to understand real-world processes, e.g., the movement of a vehicle through a chaotic medium such as the ocean, or chemical reactions it can be difficult to model the exact mechanisms occurring. Instead, we may observe the probability of certain events and incorporate these into a model where transitions are taken based on these probabilities.

An example of the latter type of model can be drawn as follows:

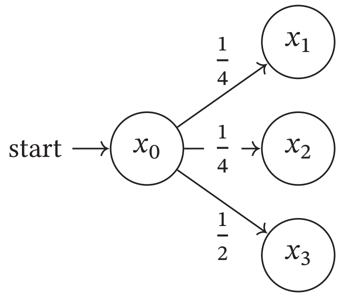

The states x_0, x_1, x_2, x_3 may each represent the position of a vehicle, or the amounts of certain species present at given times during a reaction. The movements of the vehicle and the interactions between species are then represented by the transitions which happen with the given probabilities, accounting for behaviour we are unable to model. Such models (called Markov chains in this case) thus capture a wide range of systems, whose properties can be studied together by abstracting away from details.

Another example, modelling a simplified form of the input/output interaction between a user and computer, is given by finite automata [HMU07]. The inputs are symbols/letters from an alphabet, and the system outputs whether a string of symbols is accepted or not, i.e, the computer simply outputs "yes" or "no" after each symbol. We can draw an example of such an automaton as follows:

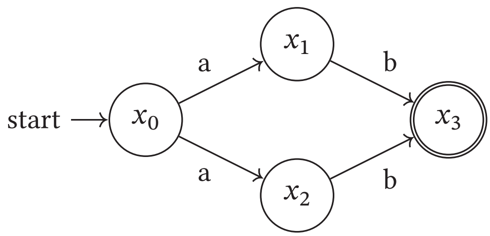

The states now represent which symbols have been read up to a certain point in the computation, with transitions being taken based on the symbol read. Computations start in the state labelled x_0. States x_0, x_1 and x_2 are not *accepting*, i.e., they output "no", while x_3 is accepting (indicated by the double outline). This means that the above automaton only accepts (says "yes") when the word "ab" is input. It can however do this in two ways: following the path via x_1 or going via x_2.

We may ask why we need this many states and two branches/paths just to accept the word "ab". In fact, this is a question we can ask more generally, and which has led to the development of the procedures of *minimisation* and *determinisation*. The first produces an automaton with the smallest number of states among those accepting the same words/strings as the original. The second produces an automaton, again accepting the same strings as the original, in which, given a symbol, there is exactly one transition possible in any state. In such an automaton, no choices need to be made to progress in the computation, so it is said to be *deterministic*.

For the above automaton, both procedures (minimisation and determinisation) will produce essentially the same automaton shown below. This is strictly speaking not the determinised automaton, as it does not show what should be done when, e.g., it reads a 'b' in state x_0. Usually, this is resolved by transitioning to a "sink state": a state which is not accepting, and in which reading any symbol loops back to the same state. We omit this state to simplify the presentation.

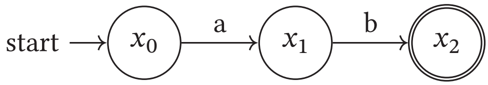

As mentioned, these procedures are made to produce automata accepting exactly the same words as the original. Indeed, the new automaton still accepts just the word "ab". We could instead require a form of equivalence based on the transition structures of the models. In the above automata, each can make a transition from x_0 with symbol 'a' to a state in which an accepting state is reached exactly when a 'b' is read (we further discuss such equivalences later in this chapter). In the above example, these definitions are equivalent, but this is not the case for all types of systems and even when they are, certain definitions may be more understandable or more suitable if we want to compute or prove equivalence of systems.

The above constructions, and their effect on various notions of equivalence raise the following questions:

1. What happens to the behaviour of state-based systems when we apply transformations to them?

2. When do we consider systems to behave in the same way, or differently, and how do we prove this?

For the case of finite automata, these questions have already been (largely) answered. Our interest in these questions comes rather from the possibility of approaching them for many types of system at once, including: (labelled) transition systems; probabilistic systems; and even arbitrary combinations of labelled, non-deterministic and probabilistic transitions. This will be facilitated by a coalgebraic view of state-based systems, which is well-established as an abstract framework for this purpose.

We extend the coalgebraic understanding of transformations and their effect on certain coalgebraic equivalences by giving new conditions ensuring that transformations preserve and/or reflect said equivalences. We further develop means for showing differences in behaviour (non-equivalence) of systems in a step-by-step (inductive) manner.

The next sections will give an outline of the coalgebraic approach to systems, including the existing notions of equivalence and their generalisation to *coinductive predicates* in *fibrations*. This general setting will be used extensively in the first part of this thesis, and has inspired the work of the second half. Along the way, we will refine the above questions to those which we aim to answer in this thesis.

Coalgebras as systems

Universal coalgebra uses the language of category theory to speak about a broad spectrum of systems involving states and transitions [Jac16]. This general "theory of systems" approach was exposited in detail by Rutten [Rut00]. Under this view, a system consists of an object X, which we think of as the collection of states, and a mapping $\gamma\colon X \to BX$, which we think of as sending these states to (structured) successors. Intuitively, successors are simply states which the system transitions to, with the form of transition being additional information. In coalgebra, this information is contained in the mapping γ, with the type of transition determined by the object BX, given by an *endofunctor* B.

As an example, consider again the first finite automaton from earlier. Its states form a set $\{x_0, x_1, x_2, x_3\}$ and transitions are made based on letters of an alphabet, $\Sigma = \{\mathrm{a}, \mathrm{b}\}$. Then, we must specify for each state and letter, the possible states transitioned to, and which states are accepting. This can be done by giving a function $\gamma\colon X \to 2 \times \mathcal{P}(X)^{\Sigma}$, which is a coalgebra for the functor $BX = 2 \times \mathcal{P}(X)^{\Sigma}$ on Set. We can see this as two maps $o\colon X \to 2$ and $t\colon X \to \mathcal{P}(X)^{\Sigma}$, specifying the output/acceptance and the transitions respectively. In the example, we would have $o(x_3) = 1$ and $o(x) = 0$ for all other states as x_3 is the only accepting state. Further, $t(x_0)(\mathrm{a}) = \{x_1, x_2\}$ and $t(x_0)(\mathrm{b}) = \varnothing$ with the other transitions defined similarly. Other examples of systems definable as coalgebras include: Mealy machines; Tree automata; Markov chains; Segala systems [Sok11]; and stream systems [Rut00; Jac16].

Together with a notion of morphism, being maps between state spaces which commute with the coalgebra structure, the coalgebras for a functor form a category. These categories will form part of the setting for most of the coming work.

Coalgebraic Equivalences

One of the aims of the field of universal coalgebra, is that once the type of system (essentially the functor B) has been chosen, many further definitions and properties should come "for free". A central example of such a property is the definition of which systems (or states) we consider to have the same behaviour.

Categorically, the canonical notion of equivalence is given by the coalgebra morphisms. Two states are considered *(coalgebraically) behaviourally equivalent* if they can be mapped by some morphism onto the same state of some coalgebra which we are free to choose.

Before the study of coalgebra, there was already the notion of *bisimulation* between processes, introduced in concurrency theory by Milner [Mil80] and Park [Par81]. Their definitions were based on the idea that states of a system are equivalent if they can copy each other's behaviour. Bisimulations were later understood as those relations preserved by transitions in the system, seen clearly in Milner's characterisation of bisimulations as post-fixed points of a monotone operator [Mil89].

This view of bisimulations as relations preserved under transitions was generalised to coalgebras by Hermida and Jacobs, using *functor lifting* [HJ98]. This assumes some functor $\overline{B}$ which, in the case of relations on sets, acts on objects by taking a relation $R \subseteq X \times X$ and producing a relation $\overline{B}(R) \subseteq BX \times BX$. Such a functor is a lifting of B in the sense that it goes from relations on X to relations on BX. A bisimulation for a coalgebra $\gamma : X \to BX$ can then be defined as a relation R with a coalgebra structure $R \to \overline{B}(R)$. Preservation of the relation by transitions in the coalgebra is ensured by requiring that if x and y are related by R, then $\gamma(x)$ and $\gamma(y)$ should be related by $\overline{B}(R)$. The precise meaning of this preservation is determined by the choice of lifting $\overline{B}$.

A further generalisation of such preserved properties are so-called coinductive invariants and coinductive predicates in fibrations [Her93; HJ98]. These are able to further capture, e.g.: semantics of modal logics by defining predicates of states satisfying a given formula; and behavioural distances by using liftings to (pseudo)metric spaces. While the full definitions use functor liftings to fibrations, we first introduce these concepts at the level of lattices, using relations and bisimulations as guiding examples.

(Co)Inductive Predicates

As mentioned above, we can view bisimulations as relations which are preserved under taking transitions in a state-based system. Such preservation can be characterised by being a post-fixed point of certain monotone operators on lattices [Par81; San98; San11]. Post-fixed points of an operator $\varphi\colon L \to L$ on a lattice L are points $p \in L$ such that $p \leq \varphi(p)$. The operator is said to be monotone if it preserves the order of the lattice. Monotonicity in the setting of *complete* lattices allows us to speak of the *greatest* (post-)fixed point.

For the case of bisimulations, the lattice is that of all relations on the state space of a given transition system. Then the transition structure allows the definition of an operator φ such that for a relation R, the relation $\varphi(R)$ consists of those pairs which we can think of as making simultaneous transitions into R. Thus, if $R \subseteq \varphi(R)$ (R is a post-fixed point of φ), then R is indeed preserved under transitions. The greatest (post-)fixed point of φ, which exists because the lattice of relations is complete, is then bisimilarity.

In the theory of coalgebra, properties beyond those definable in the setting of sets and relations are also of interest. Among these are predicates, *equivalence* relations, and (pseudo)metrics. A unifying approach to study such properties, is given by *fibrations*. At its core, a fibration consists of two categories $\mathcal{C}$ and $\mathcal{E}$ with a functor $p\colon \mathcal{E} \to \mathcal{C}$. In our context, the category $\mathcal{C}$ will be the category in which we define coalgebras, and the category $\mathcal{E}$ will be the category of spaces of properties. For example, $\mathcal{C}$ may be the category of sets and functions, and $\mathcal{E}$ the category of all relations on all sets. For each set X, the category $p^{-1}(X)$ is that of all relations $R \subseteq X \times X$ on that set X, in which we have previously defined bisimulations/bisimilarity.

A generalisation of the definition of bisimulation due to Hermida and Jacobs is then possible assuming a lifting of the coalgebraic behaviour functor $B\colon \mathcal{C} \to \mathcal{C}$ to the category $\mathcal{E}$, and using the operation of *reindexing* present in fibrations. We go into the details of fibrations in the next chapter. For the moment, we consider again the fibration of relations, in which reindexing (written f^*) along a function $f\colon X \to Y$ is given by inverse image. Hereby, we can take a relation $S \subseteq Y \times Y$, and obtain a new relation $f^*(S) = (f \times f)^{-1}(S)$ consisting of those pairs mapped by $f \times f$ into S. Given a coalgebra $\gamma\colon X \to BX$ and lifting $\overline{B}$ of B to relations, the composition $(\gamma\times\gamma)^{-1}\circ\overline{B}$, sends a relation $R \subseteq X \times X$ to those pairs of states (x, y) such that $(\gamma(x), \gamma(y)) \in \overline{B}(R)$. The post- and greatest fixed points of this monotone map $(\gamma \times \gamma)^{-1} \circ \overline{B}$ thus exactly recover the earlier definition of bisimulations/bisimilarity.

It is possible to apply this construction in any fibration for which the category $\mathcal{E}$ consists of (complete) lattices, i.e., for each object $X \in \mathcal{C}$, the space $p^{-1}(X)$ should be a complete lattice. This yields the definitions of *coinductive invariants* (post-fixed points of $\gamma^* \circ \overline{B}$)) and *coinductive predicates* (greatest fixed points of the same maps), which instantiate to the aforementioned examples of interpretations of modal formulas (see Chapter 4) and behavioural distances (see [Bal+18; BKP23; Spr+21; Kom+21] and Chapters 6 and 7).

This broad view of equivalences and other properties of coalgebras will be utilised throughout much of this thesis, and with these ideas in place, we can give a more detailed sketch of the problems we tackle in the coming chapters.

1.1 Comparing (Co)Inductive Predicates

As the example at the start of this chapter illustrated, it can be useful to transform a model of a system into a smaller one, or one in which there are no choices to be made between transitions. This second very well-known transformation of *determinisation*, was described coalgebraically in [Sil+13] and applied to such systems as: partial Mealy machines; Moore automata; and Rabin probabilistic automata. It was later shown to arise as a composition of liftings of adjoint functors to categories of coalgebras in [JSS15].

Similar constructions starting from adjunctions, whose functors are lifted to coalgebras, can be identified in the literature. Instances are: a determinisation procedure for systems combining non-determinism and probabilistic transitions [BSS21]; and ultrafilter extensions of coalgebras [KKP05]. In each case, the adjunction is a free-forgetful adjunction involving algebraic structure. The first stage of each construction is to add algebraic structure through a lifting of the free functor to obtain a system with the desired property, i.e., determinism or the coincidence of bisimilarity and modal equivalence. The second stage is then to transport this new coalgebra back to the original category by forgetting the algebraic structure through a lifting of the forgetful functor. More details of these constructions are given in Chapter 3.

For the approach in the earlier works of [Sil+13] and [JSS15], a partial answer to the question of the effect on behaviours of the constructions can be given using the assumption of a *distributive law* (see also Chapter 2). This entails that the adjunction underlying the constructions lifts to an adjunction between the categories of coalgebras. As the second stage of the construction is then given by a right adjoint, the final coalgebra (being a limit) will be preserved. In concrete settings, where the right adjoint is a forgetful functor (into Set), this tells us that behavioural equivalence is preserved and reflected.

For the cases of determinisation in [BSS21] and ultrafilter extensions in [KKP05], there is however no such distributive law. Instead, we are in the setting of *weak* distributive laws, in which the adjunction underlying the construction does not lift to coalgebras. It has been shown [GP20; Goy21] how such weak distributive laws give rise to the two-step determinisation construction of probabilistic automata. However, ultrafilter extensions have not been studied in this way, and the effect on behaviours has not been dealt with in the general setting of a weak distributive law.

We will consider this problem in the context of coinductive predicates in fibrations. Our aim is then to understand the effect of constructions obtained as liftings of adjoint functors on the properties captured by such predicates. However, it is not immediately clear how to connect the liftings to coalgebras and the fibrations underlying the properties of interest. This leads us to the question:

- What is a good setting for studying the transportation of coinductive predicates along adjunctions?

Our approach is based on adding to the picture a lifting of the adjunctions underlying the constructions to the fibrations of properties. We then wish to connect the liftings to coalgebras and fibrations to speak of preservation and reflection of coinductive predicates. The main question is then:

- Which properties of adjoint functors and their liftings to coalgebras and fibrations ensure that they preserve/reflect coinductive predicates?

The liftings of functors to coalgebras which we consider will be those induced by certain natural transformations called *steps* [RJL21]. The (weak) distributive laws used in earlier work are natural transformations of type $TS \to ST$ with T and S endofunctors or monads (necessarily on the same category). In the case of steps, we have for instance a natural transformation of type $QL \to BQ$ with B and L endofunctors on (possibly) different categories connected by the functor Q in the expected way. It has been shown by Garner [Gar20] that weak distributive laws between monads correspond (under mild conditions) to *weak liftings*, which are an instance of what we will call *invertible* steps in Chapter 3. In Chapter 4, we work with the more general assumption of *forward* and *backward* steps simply involving functors, going further beyond the setting of weak distributive laws.

As an example, a backward step will be a transformation of type $QL \to BQ$ as above, and this induces a lifting of Q going from L-coalgebras to B-coalgebras. The L-coalgebras may then be those with algebraic structure obtained in the determinisation procedure discussed earlier, and the B-coalgebras those in Set without this extra structure. The functor Q in such cases is a forgetful functor, and the lifting induced by a step of the above type realises the forgetting of algebraic structure on the coalgebras.

Our conditions for the preservation and reflection of coinductive invariants/predicates will then be based on existing conditions for the preservation and reflection of post- and greatest fixed points in lattices. While these conditions can be considered as-is at the level of lattices, we obtain more natural conditions at the level of fibrations, which combine the steps with liftings to fibrations. Our new conditions are furthermore easier to check, and are shown to entail the existing lattice-theoretic conditions. In our applications to determinisation and ultrafilter extensions, we can then show quite straightforwardly that the second stage of the constructions (the forgetting of algebraic structure) indeed preserve and reflect bisimilarity/coalgebraic behavioural equivalence.

In answering these questions at this level of generality, we obtain conditions applicable not only to generalisations of automata constructions. The steps from which these constructions arise, are also used in the semantics of coalgebraic modal logic [Mos99]. It will turn out that an instance of the transportation of coinductive predicates along adjunctions will allow the comparison of bisimilarity and logical equivalence, thereby providing conditions for the adequacy and expressivity of coalgebraic modal logics.

An alternative (dual) view of behaviours was recently studied by Geuvers and Jacobs [GJ21], called *apartness*. Via an *opposite lifting* approach developed in *op. cit.* (see also Chapter 4), we are also able to show expressivity by comparing pairs of apart states, and pairs of states for which there exists a modal formula distinguishing them. This application motivated a further investigation of apartness, which we present in the second part of this thesis, and introduce now.

1.2 Constructing Evidence of Different Behaviours

Apartness has its roots in Brouwer's *intuitionism*. Under that view, equality or equivalence of objects is not the primary means of comparison, it is rather *in*equality or *in*equivalence. This because in proving inequivalence, there is often the possibility to exhibit a difference by providing a *finite witness*.

For instance, how should we determine or prove that two real numbers are equal? It is easier to think of a procedure which shows them to be different: approximate their decimal expansions until a point is found at which they differ. A witness is then, e.g., the position at which this difference was found, or the prefix up to and including this point. These witnesses correspond to a rational strictly between the two reals.

In the study of coalgebra, the coinductive, infinite nature of the models makes proving equivalences such as bisimilarity similarly challenging. Also here, certainly for simple examples, procedures for determining differences are more straightforward. Thinking again of the example of finite automata, their behaviour is determined by their possible transitions given input sequences and the acceptance of states. To show a difference in behaviours of two states, we can thus give a word which in one system can lead to an accepting state, and in the other can not (in a deterministic system this is most straightforward as we have no choice of paths).

A coalgebraic perspective of these ideas was recently given by Geuvers and Jacobs [GJ21], based on *opposite relation lifting*. The main idea is that, given a lifting $\overline{B}$ defining a notion of bisimilarity, we can obtain a new lifting by taking a relation, complementing it, applying the lifting $\overline{B}$, and finally complementing again. It was shown in *op. cit.* that this yields the expected notions of apartness for a large class of functors, and that apartness of states can be derived inductively in these cases.

A type of functor not included in this class is the (sub)distribution functor, used to model probabilistic transitions. In fact, the opposite lifting to that of canonical relation lifting does not allow for finite witnesses of differences in general. We therefore ask:

- How can we obtain a notion of apartness for systems involving probabilistic transitions, for which apartness of states can be derived inductively?

We will use a characterisation of coalgebraic behavioural equivalence introduced by Aczel and Mendler [AM89], which can be used for the functors of interest to obtain coalgebraic behavioural equivalence as a coinductive predicate, and whose dualisation yields what we call *behavioural apartness*. For the (sub)distribution functor, the apartness notion is dual to an existing definition of bisimilarity for probabilistic systems due to Larsen and Skou [LS91]. With this approach, we are able to construct derivation systems for a recursively defined class of functors, now including the subdistribution functor. These systems will be shown to be sound and complete with respect to the definition of behavioural apartness. To ensure the usability of this system, also for systems with an infinite state space, we further show that at each proof step it suffices to restrict attention to states reachable in one step.

Apartness for Behavioural Distances

The (in)equivalences obtained from relation lifting and behavioural equivalence on probabilistic systems share the disadvantage, observed for probabilistic bisimilarity in [GJS90], that small variations in transition probabilities can change the (in)equivalence of states. One way to improve this lack of robustness, is to use *behavioural distances*, which give a measure of the similarity or difference between states. They do this by assigning to each pair of states a value, e.g., in the unit interval quantifying this similarity or difference. There are a number of definitions present in the literature for such distances, which we discuss in Chapter 6, including some possible intuitions.

Not yet discussed in the literature is a corresponding form of apartness. However, the meaning of apartness for behavioural distances is less clear than for the qualitative (two-valued) setting. Distances already give a measure of how far apart states are, so new definitions analogous to opposite relation lifting may not be needed. So, we ask:

- What is a good notion of apartness corresponding to behavioural distances?

In this case, "good" can again be considered to mean (at least) that the property can be proved inductively. We will show that this can be done for *lower bounds* on behavioural distances on labelled Markov chains: systems with probabilistic transitions and labels on the states. In our chosen ordering, such bounds will show states to be at least a certain distance apart. The derivation system we obtain will be sound and *approximately* complete with respect to the standard definition of behavioural distance for these systems [Des+99]. It is not possible in general to prove the exact distance between states as a lower bound (with finite proofs), but we can give a proof of lower bounds arbitrarily close to this *true* value. As was the case for proofs of behavioural apartness, it will suffice to consider one-step reachable states (i.e, those states in the supports of successor distributions). This again allows the application of our proof system to systems with an infinite state space.

As one of our applications in the first part of this thesis shows, distinguishing states can also be done using formulas of a modal logic. There, and in earlier work establishing expressivity of logics [DEP98; DEP02; FKP17], a correspondence between apartness (or non-bisimilarity) and distinguishing formulas has been shown. A similar correspondence is known for finite approximations of behavioural distances on labelled Markov chains and formulas in a modal logic with quantitative interpretation [RB23]. To show an equal distinguishing power of our derivation system, we give constructions going between proof trees and modal formulas exhibiting the same lower bounds. Furthermore, optimisations of the construction of derivations for lower bounds will lead to a construction of smaller distinguishing formulas than existing techniques in many examples.

1.3 Related Work

References and discussions of closely related work can be found in each of the coming chapters. Here, we give a broader comparison with some of the important alternative lines of research.

Expressivity In Chapter 4, we apply our conditions for preservation and reflection of coinductive predicates to show adequacy and expressivity of a coalgebraic modal logic. The development of such conditions is a well-studied topic, as we discuss here.

One approach focuses on properties of the modalities of the modal logic. A condition for expressivity based on separating sets of modalities was given by Pattinson [Pat04; Pat01], with existence results for such modalities given by Schröder in [Sch08]. These ideas have since been extended to quantitative modal logics in, e.g., [WS20; Gon+23; For+23].

At this quantitative level, expressivity has also been established by showing the *density* or *approximating power* of interpretations of formulas with respect to distances. This was done for so-called Kantorovich and Wasserstein liftings by König and Mika-Michalski in [KM18] and for codensity lifting by Komorida et al. in [Kom+21].

Further, there is an approach based on deriving behavioural *conformances* (such as distances) from a given logic, studied by Beohar et al. in [Beo+23; Beo+24]. A notion of *compatibility* then ensures that, e.g., the behavioural distance obtained in this way coincides with the logical one.

Our approach in the context of expressivity of coalgebraic modal logic arises from our conditions for the preservation and reflection of coinductive predicates. It is most closely related to the approach in [KR21], also based on liftings to fibrations and comparison of coinductive predicates. The lattice-theoretic conditions we use will also be related to those in [Beo+23; Beo+24], but we do not take the same approach of deriving conformances from logics: we are focussed on conditions for expressivity of given logics with respect to given conformances.

(Co)Inductive Proof Systems While we have chosen to compare behaviours of systems through apartness proofs in the second part of this thesis, others have aimed to adapt the notion of proof or computation to match the coinductive nature of systems.

One approach is circular coinduction, introduced by Roşu and Goguen [RG00] and further developed proof theoretically, and algorithmically in, e.g., [RL09; GLR00; HMS05; LR07; Luc+09]. The idea is to formalise the intuition that coinductive objects are equivalent if our attempts to distinguish them return infinitely to the same proof obligation.

Less closely related, in the sense that they do not necessarily concern proofs of equivalence, are the topics of (cyclic) proofs for logics involving (co)induction [Bro05] and computing with infinite types [Coq93]. These may use well-foundedness to show the soundness of proofs in which parts of the reasoning repeat infinitely, or require "productivity" to allow computation with infinite objects. The thesis of Basold [Bas18] deals extensively with induction and coinduction in category and type theory, and includes a detailed historical overview.

Distinguishing Formulas The use of (modal) formulas as evidence for inequivalence goes back to an algorithm by Cleaveland [Cle90] (based on the partition refinement algorithm by Kanellakis and Smolka [KS90]), which computes distinguishing formulas with respect to bisimilarity on LTSs. Korver modified this algorithm to work for branching bisimulation soon after [Kor91]. Later, Celikkan and Cleaveland [CC95] developed a procedure for giving similar "diagnostic information" for a behavioural *preorder* known as the prebisimulation preorder. Around the same time, Godskesen and Larsen gave a construction of distinguishing formulas for a type of LTS with timed behaviour [GL95].

The relation between partition refinement and the computation of distinguishing formulas can also be seen in the coalgebraic algorithm of [WMS22], which uses a coalgebraic generalisation of the algorithm by Paige and Tarjan [PT87]. There are also approaches computing distinguishing formulas via winning strategies for the spoiler in bisimulation games [KMS20]. Games and apartness have recently been connected again in [KW24; RJB24].

1.4 Outline

Chapter 2 reviews a number of concepts required in the rest of the work. These preliminaries will be used most extensively in the first part of the thesis, especially (weak) liftings and the fibrational perspective of coinductive predicates. The second part of the thesis is more self-contained, using some coalgebra and canonical relation lifting, although knowledge of fibrations and coinductive predicates can also aid the understanding of the latter chapters.

The first part of the thesis is then made up of Chapters 3 and 4 concerning preservation and reflection of bisimilarity and coinductive predicates. Chapter 3 gives our conditions for preservation and reflection of bisimilarity defined using canonical relation lifting. The conditions obtained in this chapter will mainly concern the adjoint functors underlying the motivating applications. These applications are to the settings of ultrafilter extensions for coalgebras and determinisation of probabilistic automata, for which additional details of the required invertible steps arising from the known weak distributive laws/weak liftings are also given. Chapter 3 is based on:

- Ruben Turkenburg, Clemens Kupke, Jurriaan Rot and Ezra Schoen. 'Preservation and Reflection of Bisimilarity via Invertible Steps'. In: *FoSSaCS*. 2023

Chapter 4 provides a more abstract view of preservation and reflection for coinductive predicates in fibrations. The conditions here relate to the interaction between liftings defining coinductive predicates and steps. In contrast to the previous chapter, these steps are considered separately, rather than as invertible pairs. We apply the conditions to determinisation of probabilistic automata (using lax liftings), and expressivity of coalgebraic modal logics via bisimilarity and apartness. Chapter 4 is based on:

- Ruben Turkenburg, Harsh Beohar, Clemens Kupke and Jurriaan Rot. 'Forward and Backward Steps in a Fibration'. In: *CALCO*. 2023

The second part of the thesis consists of Chapters 5 and 6, and aims to understand apartness as an inductive notion of inequivalence. Chapter 5 focuses on the development of inductive derivation systems for behavioural apartness on coalgebras; a definition obtained from dualising the existing notion of coalgebraic behavioural equivalence. A class of derivation systems is obtained for functors generated from a given grammar, which includes the subdistribution functor, not covered by earlier work using opposite relation lifting. Chapter 5 is based on:

- Ruben Turkenburg, Harsh Beohar, Clemens Kupke and Jurriaan Rot. 'Proving Behavioural Apartness'. In: *CMCS*. 2024

Chapter 6 continues the development of inductive derivation systems for apartness, now in the quantitative setting of labelled Markov chains. The focus on this system type allows a close correspondence with earlier work constructing modal formulas as witnesses for behavioural distances between states. We develop a system for proving lower bounds on these distances, which is shown to be sound and (approximately) complete with respect to a standard definition of behavioural distance. The correspondence with modal formulas with quantitative semantics is shown constructively, yielding a procedure for building witnessing formulas giving often smaller formulas than previous work. Chapter 6 is based on

- Ruben Turkenburg, Harsh Beohar, Franck van Breugel, Clemens Kupke and Jurriaan Rot. 'Constructing Witnesses for Lower Bounds on Behavioural Distances'. In: *CSL*. 2026

My Contributions

Chapters 3 to 6, are based on papers co-authored with: Harsh Beohar; Franck van Breugel; Clemens Kupke; Jurriaan Rot; and Ezra Schoen.

For each paper I was main author, responsible for most of the ideas, technical development and writing.

Further Publications

The following additional works were published during the period of my PhD, but are not part of this thesis:

- Damien Pous, Jurriaan Rot and Ruben Turkenburg. 'Corecursion Up-to via Causal Transformations'. In: *CMCS*. 2022
- Ezra Schoen, Clemens Kupke, Jurriaan Rot and Ruben Turkenburg. 'A Categorical Approach to Coalgebraic Fixpoint Logic'. In: *CMCS*. 2024
- Thorsten Wißmann, Bálint Kocsis, Jurriaan Rot and Ruben Turkenburg. 'Trees in Coalgebra from Generalized Reachability ((Co)algebraic pearl)'. In: *CALCO*. 2025

CHAPTER 2

Preliminaries

2.1 Coalgebras and Categories

We assume familiarity with the basic notions of category theory such as: categories; functors; natural transformations; (co)limits; adjoint functors.

Coalgebra For a functor $B \colon \mathcal{C} \to \mathcal{C}$, a *coalgebra* is a pair (X, γ) consisting of an object X and an arrow $\gamma \colon X \to BX$. A homomorphism from (X, γ) to (Y, ξ) is an arrow $h \colon X \to Y$ such that $\xi \circ h = Bh \circ \gamma$. Coalgebras and homomorphisms between them form a category, denoted by $\mathrm{Coalg}(B)$, or $\mathrm{Coalg}_{\mathcal{C}}(B)$ if we wish to make the underlying category explicit.

An example is coalgebras for the Set endofunctor defined on objects by $B(X) = 2 \times \mathcal{P}(X)^{\Sigma}$, used to model non-deterministic automata. A coalgebra $\gamma \colon X \to 2 \times \mathcal{P}(X)^{\Sigma}$ can be seen as two maps: $o \colon X \to 2$ specifying which states are accepting; and $t \colon X \to \mathcal{P}(X)^{\Sigma}$ giving for each state a map sending alphabet symbols to sets of successor states. Morphisms $f \colon (X, \gamma) \to (Y, \xi)$ are functions $f \colon X \to Y$ such that for all $x \in X$: $o_X(x) = o_Y(f(x))$ and for all $a \in \Sigma$ we have $f[t_X(x)(a)] = t_Y(f(x))(a)$. In other words, the coalgebra morphisms are those functions which preserve the acceptance and transitions of states.

Further examples of systems which can be captured as coalgebras and their corresponding functors, many of which we treat in more detail in the rest of this thesis, are:

- Mealy machines: $B(X) = (O \times X)^I$

- Probabilistic automata: $B(X) = \mathcal{P}(\mathcal{D}(X))^A$
- Kripke frames: $B(X) = \mathcal{P}(X)$
- Vietoris frames: $B(X) = \mathcal{V}(X)$ the Vietoris functor on CHaus or Stone
- Labelled transition systems: $B(X) = \mathcal{P}(X)^\Sigma$
- Labelled Markov chains: $B(X) = L \times \mathcal{D}(X)$

We use I, O, A, Σ and L to denote various notions of "alphabet", representing here inputs, outputs, actions, and labels.

Monads A monad is a triple $\mathcal{T} = (T, \eta, \mu)$ with: T an endofunctor on some category $\mathcal{C}$; $\eta\colon \mathsf{Id} \to T$ and $\mu\colon TT \to T$ natural transformations (called the unit and multiplication) such that the following diagrams commute:

$$\begin{tikzcd} T \arrow[r, "\eta T"] \arrow[dr, equal] & TT \arrow[d, "\mu"] & T \arrow[l, "T\eta"'] \arrow[dl, equal] \\ & T & \end{tikzcd} \qquad \begin{tikzcd} TTT \arrow[r, "T\mu"] \arrow[d, "\mu T"'] & TT \arrow[d, "\mu"] \\ TT \arrow[r, "\mu"] & T \end{tikzcd} \tag{2.1}$$

Note that we will use $\mathcal{T}$ and the underlying functor T almost interchangeably.

The following are three important examples of the above structure, which we will encounter frequently in the rest of this thesis:

- First, the powerset monad, denoted by $\mathcal{P}\colon \mathsf{Set} \to \mathsf{Set}$, given on objects by $\mathcal{P}(X) = \{S \mid S \subseteq X\}$, and on morphisms by direct image. The unit is given by $\eta(x) = \{x\}$, and multiplication given by

$$\mu(\mathcal{U}) = \bigcup_{P \in \mathcal{U}} P \tag{2.2}$$

- Second is the finitely-supported distribution monad on sets, denoted by $\mathcal{D}\colon \mathsf{Set} \to \mathsf{Set}$. It is given on objects by $\mathcal{D}(X) = \{\mu\colon X \to [0,1] \mid \sum_{x \in X} \mu(x) = 1,\ \mathrm{supp}(\mu) \text{ finite}\}$, where $\mathrm{supp}(\mu) = \{x \in X \mid \mu(x) \neq 0\}$ On functions it is given by

$$\mathcal{D}(f)(\mu)(y) = \sum\{\mu(x) \mid x \in X, f(x) = y\} \tag{2.3}$$

The unit for $\mathcal{D}$ is given by $\eta(x) = 1 \cdot |x\rangle$, and multiplication is given by

$$\mu(\nu\colon \mathcal{D}(X) \to [0,1]) = \lambda x. \sum_{p \in \mathrm{supp}(\nu)} p(x) \cdot \nu(p) \tag{2.4}$$

- Finally, we have the ultrafilter monad $\beta \colon \mathrm{Set} \to \mathrm{Set}$, which maps a set X to the set of all ultrafilters on X. A set $\mathcal{F} \subseteq \mathcal{P}X$ is a *filter* if for all $P, Q \subseteq X$, we have $P, Q \in \mathcal{F}$ iff $P \cap Q \in \mathcal{F}$. Then, we call $\mathcal{F}$ an *ultrafilter* if exactly one of the sets P and $\overline{P}$ is contained in $\mathcal{F}$ for each $P \subseteq X$. The action of β on functions is given as follows

$$\beta(f \colon X \to Y)(\mathcal{F}) = \{Q \subseteq Y \mid f^{-1}(Q) \in \mathcal{F}\} \tag{2.5}$$
$$= \{Q \subseteq Y \mid \exists P \in \mathcal{F}.\, f[P] \subseteq Q\} \tag{2.6}$$

The first equation has a possible topological intuition: the subset on Y is the largest making f "continuous". The second can be seen as taking the upward closure of the direct image, i.e., the ultrafilter generated by the direct image.

The unit for β gives for $x \in X$ the principal ultrafilter generated by x:

$$\eta(x) = \{P \subseteq X \mid x \in P\} \tag{2.7}$$

The multiplication is given by

$$\mu(\mathfrak{F} \in \beta(\beta(X))) = \{P \subseteq X \mid \mathcal{N}(P) \in \mathfrak{F}\} \tag{2.8}$$

where $\mathcal{N}(P) = \{\mathcal{F} \in \beta(X) \mid P \in \mathcal{F}\}$ (the "neighbourhoods" of P).

Algebras for a monad For a monad $\mathcal{T}$, we write $\mathcal{EM}(\mathcal{T})$ for the category of its Eilenberg-Moore algebras, and corresponding homomorphisms. This category has as objects pairs (X, α) with X an object of $\mathcal{C}$ (the category on which T is an endofunctor) and $\alpha \colon TX \to X$ such that:

$$\begin{tikzcd} X \arrow[r, "\eta_X"] \arrow[dr, equal] & TX \arrow[d, "\alpha"] \\ & X \end{tikzcd} \qquad \begin{tikzcd} TTX \arrow[r, "T\alpha"] \arrow[d, "\mu_X"'] & TX \arrow[d, "\alpha"] \\ TX \arrow[r, "\alpha"] & X \end{tikzcd} \tag{2.9}$$

Dually to coalgebras, a homomorphism of algebras $f \colon (X, \alpha) \to (Y, \beta)$ is a map $f \colon X \to Y$ such that $f \circ \alpha = \beta \circ Tf$.

For any category $\mathcal{EM}(\mathcal{T})$, there is an adjunction

$$\begin{tikzcd} \mathcal{C} \arrow[r, bend left, "\mathcal{F}"] \arrow[r, phantom, "\perp"] & \mathcal{EM}(\mathcal{T}) \arrow[l, bend left, "\mathcal{U}"] \end{tikzcd} \tag{2.10}$$

in which $\mathcal{F}(X) = (TX, \mu)$ and $\mathcal{U}(X, \alpha) = X$.

There are a number of known equivalences between Eilenberg-Moore categories (of monads on Set) and categories consisting of sets with algebraic or topological structure and structure-preserving maps. We recall three such equivalences which we make use of later. These are instances of algebraic presentations of monads [Man76].

- $\mathcal{EM}(\mathcal{P})$ is equivalent to the category of complete join semi-lattices (JSLs) and join-preserving maps [Bor94b]. Given a JSL, an algebra $\mathcal{P}(X) \to X$ can be defined by simply taking joins.

- $\mathcal{EM}(\mathcal{D})$ is equivalent to the category of convex algebras and convex maps. The convex combinations given by the operations of a convex algebra can be seen as distributions, thus defining a map $\mathcal{D}(X) \to X$. See [Dob06].

- $\mathcal{EM}(\beta)$ is equivalent to the category of compact Hausdorff topological spaces and continuous maps. The construction of an algebra $\beta X \to X$ from a compact Hausdorff space uses a notion of ultrafilter convergence. Namely, an ultrafilter $\mathcal{F} \in \beta X$ is said to converge to a point $x \in X$ if all open neighbourhoods of x are contained in $\mathcal{F}$. As Manes showed [Man69], for compact Hausdorff spaces, each ultrafilter converges to exactly one point, thus defining a map $\beta X \to X$.

2.2 (Weak) Distributive Laws and (Weak) Liftings

In a number of applications of our results in Chapters 3 and 4, we will use the definitions of (weak) liftings of monads to Eilenberg-Moore categories. Also important in those settings is the correspondence between (weak) liftings and (weak) distributive laws. We briefly recall the necessary definitions and the results showing these correspondences.

Definition 2.2.1: For monads $\mathcal{T} = (T, \eta^T, \mu^T)$ and $\mathcal{S} = (S, \eta^S, \mu^S)$, a distributive law of $\mathcal{T}$ over $\mathcal{S}$ is a natural transformation $\lambda\colon TS \to ST$ satisfying the following axioms:

$$\begin{tikzcd} & S \arrow[dl, "\eta^T S"'] \arrow[dr, "S\eta^T"] & \\ TS \arrow[rr, "\lambda"] & & ST \end{tikzcd} \tag{2.11}$$

$$\begin{tikzcd} & T \arrow[dl, "T\eta^S"'] \arrow[dr, "\eta^S T"] & \\ TS \arrow[rr, "\lambda"] & & ST \end{tikzcd} \tag{2.12}$$

$$\begin{tikzcd} TTS \arrow[r, "T\lambda"] \arrow[d, "\mu^T S"'] & TST \arrow[r, "\lambda T"] & STT \arrow[d, "S\mu^T"] \\ TS \arrow[rr, "\lambda"] & & ST \end{tikzcd} \tag{2.13}$$

$$\begin{tikzcd} TSS \arrow[r, "\lambda S"] \arrow[d, "T\mu^S"'] & STS \arrow[r, "S\lambda"] & SST \arrow[d, "\mu^S T"] \\ TS \arrow[rr, "\lambda"] & & ST \end{tikzcd} \tag{2.14}$$

◂

As shown by Beck [Bec06], these distributive laws correspond to the following notion of lifting of monads to Eilenberg-Moore categories.

Definition 2.2.2: A lifting of a monad $\mathcal{S}$ to $\mathcal{EM}(\mathcal{T})$ is a monad

$$(\widetilde{\mathcal{S}}, \eta^{\widetilde{\mathcal{S}}}, \mu^{\widetilde{\mathcal{S}}}) : \ \mathcal{EM}(\mathcal{T}) \to \mathcal{EM}(\mathcal{T})$$

such that $U\widetilde{\mathcal{S}} = \mathcal{S}U$, $U\eta^{\widetilde{\mathcal{S}}} = \eta^S U$ and $U\mu^{\widetilde{\mathcal{S}}} = \mu^S U$. ◂

Proposition 2.2.3 ([Bec06] [Gar20, Prop. 6]): There is a bijective correspondence between distributive laws of $\mathcal{T}$ over $\mathcal{S}$ and liftings of $\mathcal{S}$ to $\mathcal{EM}(\mathcal{T})$. ◂

Also of interest, especially in our applications, are the following two weaker notions, introduced in this form by Garner [Gar20] based on earlier notions of Street [Str09] and Böhm [Böh10].

Definition 2.2.4: For monads $\mathcal{T} = (T, \eta^T, \mu^T)$ and $\mathcal{S} = (S, \eta^S, \mu^S)$, a *weak* distributive law of $\mathcal{T}$ over $\mathcal{S}$ is a natural transformation $\lambda : \ TS \to ST$ satisfying the axioms of Eqs. (2.12) to (2.14). I.e., it satisfies all axioms of a distributive law except the one involving the unit of $\mathcal{T}$. ◂

Definition 2.2.5: For a monad $\mathcal{T} = (T, \eta^T, \mu^T)$ and a functor S, a weak distributive law of $\mathcal{T}$ over the functor S is a natural transformation $\lambda : \ TS \to ST$ satisfying only the axioms of Eqs. (2.11) and (2.13). ◂

The above distributive laws will give a way to obtain *weak* liftings of monads or functors to categories of algebras for a monad, satisfying weaker conditions than the equality $U\widetilde{\mathcal{S}} = \mathcal{S}U$ of Definition 2.2.2.

Definition 2.2.6: A *weak lifting of a monad $\mathcal{S}$ to $\mathcal{EM}(\mathcal{T})$* is a monad $\widetilde{\mathcal{S}} = (\tilde{S}, \eta^{\tilde{S}}, \mu^{\tilde{S}}) : \ \mathcal{EM}(\mathcal{T}) \to \mathcal{EM}(\mathcal{T})$ such that there are natural transformations:

$$U\tilde{S} \xrightarrow{\iota} SU \xrightarrow{\delta} U\tilde{S} \tag{2.15}$$

satisfying $\delta \circ \iota = \mathrm{id}_{U\tilde{S}}$ and making the following diagrams commute

$$\begin{tikzcd} & U \arrow[dl, "U\eta^S"'] \arrow[d, "\eta^S U"] \arrow[dr, "U\eta^S"] & \\ U\tilde{S} \arrow[r, "\iota"] & SU \arrow[r, "\delta"] & U\tilde{S} \end{tikzcd} \tag{2.16}$$

$$\begin{tikzcd} U\tilde{S}\tilde{S} \arrow[r, "\iota\tilde{S}"] \arrow[d, "U\mu^{\tilde{S}}"'] & SU\tilde{S} \arrow[r, "S\iota"] & SSU \arrow[r, "S\delta"] \arrow[d, "\mu^S U"] & SU\tilde{S} \arrow[r, "\delta\tilde{S}"] & U\tilde{S}\tilde{S} \arrow[d, "U\mu^{\tilde{S}}"] \\ U\tilde{S} \arrow[rr, "\iota"] & & SU \arrow[rr, "\delta"] & & U\tilde{S} \end{tikzcd} \tag{2.17}$$

◀

Definition 2.2.7: A *weak lifting of a functor S to $\mathcal{EM}(\mathcal{T})$*, is a functor $\tilde{S} \colon \mathcal{EM}(\mathcal{T}) \to \mathcal{EM}(\mathcal{T})$ and natural transformations

$$U\tilde{S} \xrightarrow{\iota} SU \xrightarrow{\delta} U\tilde{S} \tag{2.18}$$

such that $\delta \circ \iota = \mathrm{id}$. ◀

We use the same notation for all liftings, but we deal mainly with weak liftings and will make the type clear whenever confusion could arise.

Now, we recall the following two (very similar) results relating the weak variants of distributive laws and weak liftings analogously to Proposition 2.2.3; note the difference between $\mathcal{S}$ and S.

Proposition 2.2.8 ([Gar20, Prop 13]): If idempotents in $\mathcal{C}$ split, then there is a bijective correspondence between weak distributive laws of a monad $\mathcal{T}$ over a monad $\mathcal{S}$ and weak liftings of $\mathcal{S}$ to $\mathcal{EM}(\mathcal{T})$. ◀

Proposition 2.2.9 ([Goy21, Prop 4.3]): If idempotents in $\mathcal{C}$ split, then there is a bijective correspondence between weak distributive laws of a monad $\mathcal{T}$ over a functor S and weak liftings of S to $\mathcal{EM}(\mathcal{T})$. ◀

2.3 Factorisation systems

To define a standard technique for lifting functors to predicates and relations (in the next section), we first require the definition of factorisation systems. Some intuition can be taken from the case of functions on sets, and their notion of *image*. We can decompose a function $f \colon X \to Y$ into a surjection $s \colon X \to \mathrm{Im}(f)$ onto the elements of Y reached by f, and an injection $i \colon \mathrm{Im}(f) \to Y$ of these elements into the original codomain. The coming definitions generalise this idea.

In a category $\mathcal{C}$, morphisms $e\colon X \to Y, m\colon Z \to W$ are called *orthogonal* (written $e \perp m$) if for any morphisms $u\colon X \to Z, v\colon Y \to W$ such that $v \circ e = m \circ u$, there is a unique morphism $d\colon Y \to Z$ with $d \circ e = u$ and $m \circ d = v$. In a diagram

$$\begin{tikzcd} X \arrow[r, "e"] \arrow[d, "u"'] & Y \arrow[d, "v"] \arrow[dl, dashed, "d"'] \\ Z \arrow[r, "m"'] & W \end{tikzcd} \tag{2.19}$$

We may also say that e, m have the *diagonal fill-in property*.

Definition 2.3.1: A *factorisation system* on $\mathcal{C}$ is a pair $(\mathcal{E}, \mathcal{M})$ with $\mathcal{E}$ and $\mathcal{M}$ collections of morphisms of $\mathcal{C}$ such that:

1. all isomorphisms of $\mathcal{C}$ are in both $\mathcal{E}$ and $\mathcal{M}$;
2. $\mathcal{E}$ and $\mathcal{M}$ are closed under composition;
3. for any $e \in \mathcal{E}$ and $m \in \mathcal{M}$ we have $e \perp m$;
4. every arrow f of $\mathcal{C}$ can be factorised as $f = m \circ e$ with $e \in \mathcal{E}$ and $m \in \mathcal{M}$. ◂

Remark 2.3.2: A note on notation: we use $\twoheadrightarrow$ for abstract epis and $\rightarrowtail$ for abstract monos, i.e., maps in $\mathcal{E}$ and $\mathcal{M}$ respectively. ◂

When using factorisation systems to define predicate and relation liftings, it is useful to make additional restrictions, to ensure that operations such as intersections interact well with the factorisations. We use *logical* factorisation systems, as introduced by Jacobs in [Jac16, Def. 4.3.2], and recalled here.

Definition 2.3.3: A factorisation system $(\mathcal{E}, \mathcal{M})$ on a category $\mathcal{C}$ is *logical* if

1. $\mathcal{M} \subseteq \mathsf{Mono}$;
2. $\mathcal{M}$ is closed under pullback;
3. for any $e \in \mathcal{E}$ and $m \in \mathcal{M}$, the pullback of e along m is again in $\mathcal{E}$. ◂

To recover the earlier example of image factorisation in Set, we can take $\mathcal{E}$ to be the surjective functions and $\mathcal{M}$ to be the injective functions.

Another important example is that of *regular categories*, where maps factorise as a *regular* epi (i.e., an epi which is the coequaliser of some parallel pair of morphisms) followed by a mono. In fact, for a finitely complete category, having such factorisations which are furthermore stable under pullbacks, is equivalent to being regular [Gra21, Thm. 1.14].

2.4 Predicates, Relations, and Liftings

Spaces of predicates and relations on the state spaces of coalgebras will play a large role in the coming work. They will form the categories where we define such objects as behavioural equivalences and interpretations of logical formulas.

To make this general enough for our purposes, we start with a category $\mathcal{C}$ which is complete and well-powered, meaning subobjects of a given object form a set. Fixing a subset $\mathcal{M} \subseteq \mathsf{Mono}$ of monomorphisms, we define an ordering on $\mathcal{M}$ in the usual way. Namely, for $m_1 \colon X \rightarrowtail Z$ and $m_2 \colon Y \rightarrowtail Z$, we define $m_1 \leq m_2$ if there exists an arrow $l \colon X \to Y$ such that $m_2 \circ l = m_1$. This ordering induces an equivalence relation $\leq \cap \geq$ on $\mathcal{M}$, and by *subobject* we will mean an equivalence class of this relation. We can thus speak of the subobject represented by some $m \in \mathcal{M}$. The standard definition of subobject is recovered by taking $\mathcal{M} = \mathsf{Mono}$.

The category $\mathsf{Pred}(\mathcal{C})$ of predicates then has as objects such subobjects represented by maps $m \colon S \rightarrowtail X$ in $\mathcal{M}$. A morphism from $m \colon S \rightarrowtail X$ to $n \colon T \rightarrowtail Y$ is a map $u \colon X \to Y$ in $\mathcal{C}$ so that there exists a (necessarily) unique map making the following diagram commute:

$$\begin{tikzcd} S \arrow[r, dashed] \arrow[d, tail] & T \arrow[d, tail] \\ X \arrow[r, "u"] & Y \end{tikzcd} \tag{2.20}$$

For a complete category $\mathcal{C}$, we can also form the category $\mathsf{Rel}(\mathcal{C})$ of (homogeneous) relations, with objects being subobjects $r \colon R \rightarrowtail X \times X$ of the product of an object $X \in \mathcal{C}$ with itself. Arrows are then just as for $\mathsf{Pred}(\mathcal{C})$.

For Set with $\mathcal{M} = \mathsf{Mono}$, $\mathsf{Rel}(\mathsf{Set})$ consists of subsets of the binary product of underlying sets as usual, and maps between relations constitute maps between the products sending R to S, i.e., $x\,R\,y$ implies $u(x)\,S\,u(y)$. Objects of $\mathsf{Rel}(\mathsf{Stone})$ are closed relations, as the image of a mono representing a subobject is homeomorphic (related by a continuous bijection) to its domain, and images of continuous functions are compact and thus closed. In the case of an Eilenberg-Moore category for a monad T, objects of $\mathsf{Rel}(\mathcal{EM}(T))$ are congruences, as the map into the product is an algebra morphism.

We can now define the (canonical) predicate and relation liftings $\mathsf{Pred}(F)$ and $\mathsf{Rel}(F)$ of a functor $F \colon \mathcal{C} \to \mathcal{D}$ via the following factorisations, where $p \colon P \rightarrowtail X$ is a predicate and $r \colon R \rightarrowtail X \times X$ a relation:

$$\begin{tikzcd} FP \arrow[rr, "Fp"] \arrow[dr, two heads, "e"'] & & FX \\ & \mathsf{Pred}(F)(P) \arrow[ur, tail, "m"'] & \end{tikzcd} \tag{2.21}$$

(2.22)
$$\begin{tikzcd} FR \arrow[r, "Fr"] \arrow[d, hook] & F(X\times X) \arrow[rr, "\langle F\pi_1, F\pi_2\rangle"] & & FX \times FX \\ {} \arrow[r, two heads, "e"] & \mathsf{Rel}(F)(R) \arrow[rr, tail, "m"] & & {} \arrow[u] \end{tikzcd}$$

This defines functors $\mathsf{Pred}(F)\colon \mathsf{Pred}(\mathcal{C}) \to \mathsf{Pred}(\mathcal{D})$ and $\mathsf{Rel}(F)\colon \mathsf{Rel}(\mathcal{C}) \to \mathsf{Rel}(\mathcal{D})$ respectively, with the actions on arrows defined by orthogonality.

2.5 Fibrations

We give the basic definitions related to fibrations (for details see [Jac01]).

Given a functor $p\colon \mathcal{E} \to \mathcal{C}$, a morphism $b\colon R \to S$ in $\mathcal{E}$ is *(p-)Cartesian* over $f\colon X \to Y$ in $\mathcal{C}$, if $p(b) = f$ and for every $c\colon T \to S$ s.t. $p(c) = f \circ g$ for some $g\colon p(T) \to X$, there is a unique $d\colon T \to R$ with $c = b \circ d$. In a diagram:

(2.23)
$$\begin{tikzcd} T \arrow[dr, dashed, "\exists! d"'] \arrow[drr, bend left, "\forall c"] & & \\ & R \arrow[r, "b"'] & S \end{tikzcd}$$
$$\begin{tikzcd} pT \arrow[dr, "\forall g"'] \arrow[drr, bend left, "pc = f\circ g"] & & \\ & X \arrow[r, "f = pb"'] & Y \end{tikzcd}$$

A functor $p\colon \mathcal{E} \to \mathcal{C}$ is now a (Grothendieck) fibration if for all objects $S \in \mathcal{E}$ and arrows $f\colon X \to p(S)$, there is a Cartesian arrow $b\colon R \to S$ in $\mathcal{E}$ with $p(b) = f$ (and thus also $pR = X$). We say that R is *above* pR and $b\colon R \to S$ is *above* $p(b)\colon p(R) \to p(S)$. We will also call $\mathcal{C}$ the base category and $\mathcal{E}$ the total category of the fibration.

For an object $X \in \mathcal{C}$, the *fibre above* X is the subcategory $\mathcal{E}_X$ of $\mathcal{E}$ whose objects are those objects in $\mathcal{E}$ above X, and arrows are above the identity on X. Choosing for every $S \in \mathcal{E}$ and $f\colon X \to p(S)$ in $\mathcal{C}$ a Cartesian lifting, denoted $\overline{f}(S)\colon f^*(S) \to S$, yields what is called a *cleavage*. Given a cleavage, each arrow $f\colon X \to Y$ in $\mathcal{C}$ gives rise to a reindexing functor $f*\colon \mathcal{E}_Y \to \mathcal{E}_X$. On objects, this maps an $S \in \mathcal{E}_Y$ to the domain $f^*(S)$ of the chosen Cartesian lifting $\overline{f}$. For arrows, it takes some $b\colon S_1 \to S_2$ in $\mathcal{E}_Y$ and yields the map obtained from the Cartesian property of $\overline{f}(S_2)$.

We assume below that reindexing functors have left adjoints $\coprod_f \dashv f^*$ (called *direct image*). This is equivalent to the condition that both $p\colon \mathcal{E} \to \mathcal{C}$ and $p^{\mathrm{op}}\colon \mathcal{E}^{\mathrm{op}} \to \mathcal{C}^{\mathrm{op}}$ are fibrations, in which case, p is also called a *bifibration*.

Given fibrations $p\colon \mathcal{E} \to \mathcal{C}$ and $q\colon \mathcal{F} \to \mathcal{D}$, a *morphism of fibrations* from p to q is a pair of functors $(\overline{F}\colon \mathcal{E} \to \mathcal{F}, F\colon \mathcal{C} \to \mathcal{D})$ such that $q \circ \overline{F} = F \circ p$. In that case, for every object we have a restriction $\overline{F}_X\colon \mathcal{E}_X \to \mathcal{F}_{FX}$, denoted by $\overline{F}$ if the type is clear from the context. We will also call $\overline{F}$ a lifting of F. If $\overline{F}$ preserves Cartesian morphisms, it is called *fibred*. In a fibration where the fibres are posets, this is equivalent to having the equation $\overline{F}_X \circ f^* = (Ff)^* \circ \overline{F}_Y$ for all morphisms $f\colon X \to Y$.

2.5.1 Example: Predicates and Relations

A recurring class of examples in this work are the fibrations of predicates and relations. Given a complete, well-powered category $\mathcal{C}$, the objects of the total categories $\mathsf{Pred}(\mathcal{C})$ and $\mathsf{Rel}(\mathcal{C})$ are predicates and relations as defined in Section 2.4. Then the functors $p\colon \mathsf{Pred}(\mathcal{C}) \to \mathcal{C}$ and $q\colon \mathsf{Rel}(\mathcal{C}) \to \mathcal{C}$ mapping a predicate $P \rightarrowtail X$ to X and a relation $R \rightarrowtail X \times X$ to X respectively, are fibrations over $\mathcal{C}$.

In these fibrations, reindexing can be computed by taking the pullback of a predicate or relation along the relevant map. For example for a predicate $P \rightarrowtail Y$ and map $f\colon X \to Y$, the reindexing $f^*(P)$ can be obtained as follows:

$$\begin{tikzcd} f^*(P) \arrow[r] \arrow[d, tail, "f^*(p)"'] \arrow[dr, phantom, "\lrcorner", very near start] & P \arrow[d, tail, "p"] \\ X \arrow[r, "f"] & Y \end{tikzcd} \tag{2.24}$$

For a relation, we simply take the pullback along a map $f \times f$.

In case the underlying category $\mathcal{C}$ furthermore has a logical factorisation system $(\mathcal{E}, \mathcal{M})$ and subobjects in $\mathcal{C}$ are defined with respect to the set $\mathcal{M}$, then the left adjoint $\coprod_f$ to reindexing f^* in $\mathsf{Pred}(\mathcal{C})$ is obtainable as the following factorisation:

$$\begin{tikzcd} S \arrow[r, two heads] \arrow[d, tail] & \coprod_f(S) \arrow[d, tail] \\ X \arrow[r, "f"] & Y \end{tikzcd} \tag{2.25}$$

and similarly for $\mathsf{Rel}(\mathcal{C})$. See [Jac16, Prop. 4.3.5] for further details.

Example 2.5.1: In Set, subobjects are just subsets $U \subseteq X$, and reindexing is inverse image, i.e., $f^*(V \subseteq Y) = \{x \in X \mid f(x) \in V\}$. Similarly, relations are subsets $R \subseteq X \times X$, with $f^*(S \subseteq Y \times Y) = \{(x_1, x_2) \in X \times X \mid (f(x_1), f(x_2)) \in S\}$. Notice Set is complete, and the collection of subsets of a set is its powerset which is again a set. Each powerset is thus a complete lattice with join and meet given by union and intersection, respectively.

For a monad $T\colon \mathsf{Set} \to \mathsf{Set}$, let $\mathcal{EM}(T)$ be the category of Eilenberg-Moore algebras. Then, the category $\mathsf{Pred}(\mathcal{EM}(T))$, consists of subalgebras, and $\mathsf{Rel}(\mathcal{EM}(T))$ consists of congruences, that is relations that are closed under applications of the algebra structure (not necessarily equivalence relations). ◂

Predicate and Relation liftings

Recall the definitions of predicate and relation liftings from Section 2.4. In the context of a functor $F\colon \mathcal{C} \to \mathcal{D}$ between complete well-powered categories with logical factorisation systems, these are liftings to the predicate and relation fibrations in the sense that they commute with the functors underlying the fibrations:

$$\begin{tikzcd} \mathsf{Pred}(\mathcal{C}) \arrow[r, "\mathsf{Pred}(F)"] \arrow[d] & \mathsf{Pred}(\mathcal{D}) \arrow[d] \\ \mathcal{C} \arrow[r, "F"] & \mathcal{D} \end{tikzcd} \qquad \begin{tikzcd} \mathsf{Rel}(\mathcal{C}) \arrow[r, "\mathsf{Rel}(F)"] \arrow[d] & \mathsf{Rel}(\mathcal{D}) \arrow[d] \\ \mathcal{C} \arrow[r, "F"] & \mathcal{D} \end{tikzcd} \tag{2.26}$$

2.5.2 $\mathsf{CLat}_\sqcap$ fibrations

We will apply the theory of fibrations to capture properties of (co)algebras as so-called *(co)inductive invariants/predicates*. In this context, it is common to assume a fibration $p\colon \mathcal{E} \to \mathcal{C}$ for which all fibres $\mathcal{E}_X$ for objects $X \in \mathcal{C}$ are complete lattices and reindexing preserves the meet operation in these lattices. We will call such a fibration a $\mathsf{CLat}_\sqcap$-fibration. Note that these may also be called topological functors (see, e.g., [AHS09]).

The assumption ensures that the fibrations have many desirable properties, while being general enough for our purposes of defining coinductive predicates. We will make this assumption explicit where necessary. In particular, the assumption of completeness of the fibres ensures that the coinductive predicates we present soon will exist, as they are defined as (greatest) fixed points. Such fibrations are also always bifibrations, so that each reindexing functor f^* has a left adjoint $\coprod_f$, as discussed earlier. For a more detailed treatment of $\mathsf{CLat}_\sqcap$-fibrations see, e.g., [Kom+19; Spr+21].

Subobjects and Relations As can be seen in our earlier examples, we will be interested in the fibrations of predicates and relations in some base category, and the lifting of functors to these total categories.

Assuming the base category is complete and well-powered, the fibres $\mathsf{Pred}(\mathcal{C})_X$ are complete lattices [Bor94a, Cor. 4.2.5]. Since reindexing is defined by pullback, it preserves meets, so that p is a $\mathsf{CLat}_\sqcap$-fibration. Objects of $\mathsf{Rel}(\mathcal{C})$ are a special case of predicates, so that this is also a $\mathsf{CLat}_\sqcap$-fibration.

We will later also wish to restrict attention to *equivalence* relations. We can define the subcategory $\mathsf{ERel}(\mathcal{C})$ of equivalence relations, defined internally to the category $\mathcal{C}$ as follows (see also, e.g., [Bor94b, Prop. 2.5.5]):

Definition 2.5.2: A relation $r\colon R \rightarrowtail X \times X$ in a category $\mathcal{C}$ is an (internal) equivalence relation if:

1. the diagonal $\Delta_X \rightarrowtail X \times X$ factors through r;
2. the map $s \circ r$ factors through r, where s is the (unique) map $X \times X \to X \times X$ such that $\pi_1 \circ s = \pi_2$ and vice versa;
3. the tuple map $\langle \pi_1 \circ r \circ \rho_1, \pi_2 \circ r \circ \rho_2 \rangle$ factors through r where ρ_1, ρ_2 are as in the following pullback:

$$\begin{tikzcd} R \times_X R \arrow[r, "\rho_2"] \arrow[d, "\rho_1"'] & R \arrow[d, "\pi_1 \circ r"] \\ R \arrow[r, "\pi_2 \circ r"'] & X \end{tikzcd} \tag{2.27}$$ ◂

Reindexing and meets for such equivalence relations are defined in the same way as for relations, i.e., as pullbacks. This turns $\mathsf{ERel}(\mathcal{C}) \to \mathcal{C}$ again into a $\mathsf{CLat}_\sqcap$-fibration.

2.5.3 Invariants and Coinductive Predicates

We will now recall the notion of coinductive invariants and predicates, defined as post and greatest fixed points of certain monotone maps respectively (see also [Jac16; HKC18]). Working in a $\mathsf{CLat}_\sqcap$-fibration ensures that these monotone maps always have such fixpoints.

Assuming a fibration $p\colon \mathcal{E} \to \mathcal{C}$, a coalgebra $\gamma\colon X \to BX$ with X in $\mathcal{C}$, and a lifting $\overline{B}\colon \mathcal{E} \to \mathcal{E}$ of $B\colon \mathcal{C} \to \mathcal{C}$, we define the monotone map $\gamma^* \circ \overline{B}_X : \mathcal{E}_X \to \mathcal{E}_X$. Then a post-fixed point $R \leq \gamma^* \circ \overline{B}_X(R)$ will also be called an *invariant*, and the greatest fixed point $\nu(\gamma^* \circ \overline{B}_X)$ will be called the *coinductive predicate* defined by $\overline{B}$ on γ. Instantiating this to the category Set, and the fibration with Rel(Set) as total category, we can consider the canonical relation lifting (2.22) of B, given explicitly by

$$\mathsf{Rel}(B)(R) = \{(y_1, y_2) \in BX \times BX \mid \exists z \in BR. B\pi_1(z) = y_1 \\ \wedge B\pi_2(z) = y_2\} \tag{2.28}$$

As mentioned earlier, reindexing for the relation fibration is given by pullbacks, which amounts to taking the inverse image, so that:

$$\gamma^* \circ \mathsf{Rel}(B)(R) = \{(x_1, x_2) \in X \times X \mid \exists z \in BR. B\pi_1(z) = \gamma(x_1) \\ \wedge B\pi_2(z) = \gamma(x_2)\} \tag{2.29}$$

A post-fixed point $R \leq \gamma^* \circ \mathsf{Rel}(B)(R)$ of such a monotone map recovers the usual notion of *coalgebraic bisimulation.* The greatest fixed point $\nu(\gamma^* \circ \mathsf{Rel}(B))$ is then called bisimilarity.

These notions cover many more examples than bisimulations and bisimilarity. For a simple instance, take B to be the powerset functor $\mathcal{P}\colon \mathsf{Set} \to \mathsf{Set}$, and $\overline{\mathcal{P}}\colon \mathsf{Pred} \to \mathsf{Pred}$ with $\overline{B}(P \subseteq X) = \{S \subseteq X \mid P \cap S \neq \emptyset\}$. A $\mathcal{P}$-coalgebra is a transition system, and the coinductive predicate on it defined by $\overline{\mathcal{P}}$ is the set of all states which have an infinite path. For other examples of coinductive predicates defined in this way see, e.g., [HKC18; BKP23; Spr+21; HJ04; Kom+19].

Part I

Comparing (Co)Inductive Predicates

CHAPTER 3

Preservation and Reflection of Bisimilarity via Invertible Steps

Just a monoid, that's crazy
about working! (5)

3.1 Introduction

DISTRIBUTIVE LAWS between a monad T and a functor B are ubiquitous in the theory of coalgebras. They capture various forms of interaction between algebras and coalgebras, including: structural operational semantics [TP97; Kli11]; efficient proof techniques [Bon+17]; and a general coalgebraic determinisation procedure which applies to a wide range of automata and other state-based systems [Sil+13; BMS13; JSS15].

As shown in [JSS15], the central categorical idea of this general determinisation procedure is to transport a BT-coalgebra to a coalgebra for the lifting $\hat{B}$ of B to the Eilenberg-Moore category $\mathcal{EM}(T)$ induced by the distributive law. Behavioural equivalence in $\mathcal{EM}(T)$ then amounts to desired notions of equivalence. Examples of systems which have been studied in this way include: language equivalence of non-deterministic automata; weighted automata [Bon+12]; Mealy and Moore machines with side effects [Sil+13]; and various types of trace equivalence of transition systems [Bon+16].

An illustrative *non-example* of this general determinisation procedure is in a natural construction of belief-state transformers from probabilistic automata (PA), which feature both non-determinism and probabilities. This construction moves the probabilistic branching in a PA from the transitions into convex structure on the states, generalising the well-known powerset construction which moves non-deterministic branching into join semi-lattice structure.

From a categorical perspective, the problem in the case of probabilistic automata is related to the classical result that there is no suitable distributive law of the probability distribution monad $\mathcal{D}$ over the powerset monad $\mathcal{P}$ [Var03] (also see [ZM22; KS18] for other non-existence results of distributive laws). Hence, it seems that general determinisation via distributive laws is not applicable here.

Nevertheless, in [BSS21] a coalgebraic account of the construction of belief-state transformers is given as a two-stage process:

1. from probabilistic automata to coalgebras in $\mathcal{EM}(\mathcal{D})$, which are a type of labelled transition systems over convex algebras;

2. from these coalgebras in $\mathcal{EM}(\mathcal{D})$ back to plain transition systems in Set, yielding the belief-state transformer.

A key result in *op. cit.* is that the second stage preserves and reflects behavioural equivalence, defined there by the existence of a kernel bisimulation [Sta11; Kur00; Wol00]. This shows that behavioural equivalence of the coalgebras in $\mathcal{EM}(\mathcal{D})$ coincides with behavioural equivalence on the labelled transition systems obtained by forgetting the convex structure. This further coincides with what is called distribution bisimilarity in [BSS21], i.e., bisimilarity on labelled transition systems (LTSs) whose states are distributions.

In [BSS21; Goy21] it was shown that the two-stage construction above in fact arises from a canonical *weak distributive law* of $\mathcal{D}$ over $\mathcal{P}$ [GP20]. Weak distributive laws of the form used in *op. cit.* correspond to so-called weak liftings of monads/functors to Eilenberg-Moore categories [Gar20] which, as shown in [GP20], yield a new generalised determinisation procedure which covers the example of probabilistic automata, and precisely instantiates to the two stages above. Further examples are the treatment of alternating automata via weak distributive laws in [GPA21], and weak distributive laws for combining non-determinism with semimodules in [BS21].

However, the result for probabilistic automata that the second stage above preserves and reflects behavioural equivalence has not yet been accounted for in the abstract theory of determinisation via weak distributive laws.

In this chapter we provide such an account, starting from a more general setting than weak distributive laws: what we call *invertible steps*. While weak liftings by definition apply to Eilenberg-Moore categories of monads, (invertible) steps can be defined in the context of arbitrary adjunctions. A *step*, named as such in [RJL21], is a natural transformation of a certain type allowing the lifting of the left adjoint to coalgebras. Such steps are a widely occurring phenomenon, for instance in the semantics of coalgebraic modal logic, testing semantics and trace semantics (see [RJL21] for an overview). To further generalise the second stage of the aforementioned constructions, the key idea is to assume a natural transformation in the opposite direction to the existing notion of step, which yields a lifting of the right adjoint. We also assume that this new transformation is a right inverse to a step, together forming what we will call an *invertible step*.

In this setting of an invertible step, we are able to show that the second stage of the two-stage construction preserves and reflects bisimilarity, under mild conditions. As a consequence, we recover the above-mentioned results on preservation and reflection of behavioural equivalence for probabilistic automata [BSS21] for free from the abstract theory.[1] Another motivating example is that of coalgebras for the Vietoris functor on the category of Stone spaces: we obtain that bisimilarity is preserved and reflected by (a lifting of) the forgetful functor from Stone spaces to sets, recovering the main result in [BFV10].

This is relevant in a coalgebraic presentation [KKP05] of ultrafilter extensions, a standard construction in modal logic [BRV01]. In this coalgebraic form, the construction has a similar structure to the generalised determinisation procedures, and it is again important that forgetting topological structure does not change our notion of equivalence. This will be one of our running examples, and shows the benefit of the general approach based on (invertible) steps, as it does not fit directly into the setting of weak liftings, due to the category of Stone spaces (for the duality with Boolean algebras) which is not monadic over Set. However, if we move from Stone spaces to compact Hausdorff spaces, then the relevant weak lifting (or invertible step) arises precisely from the weak distributive law constructed by Garner [Gar20], and can then be restricted back to Stone spaces. The weak distributive law in *op. cit.* thus gives rise to ultrafilter extensions in modal logic.

[1] We focus on bisimilarity, but our setting allows for an easy argument that this coincides with coalgebraic behavioural equivalence defined by identification under some homomorphism in this and many related examples.

Contributions The first main contribution of this chapter is the identification of invertible steps as structure underlying existing coalgebraic formulations of automata constructions. We then show under what conditions on these steps and their underlying functors, the induced liftings to coalgebras preserve or reflect bisimulations/bisimilarity.

The first of these results provides a condition for preservation of bisimulations and bisimilarity by liftings of right adjoint functors. This is quite a straightforward condition: that the right adjoint preserves abstract epimorphisms of assumed factorisation systems. This is shown in Theorem 3.3.11.

That result uses liftings of the adjoint functors (not the whole adjunction) to the spaces of relations in which bisimulations and bisimilarity are defined. We show that, at the level of lattices, the relation lifting of the right adjoint has a left adjoint, which always preserves bisimulations (Proposition 3.3.13).

We further have two extensive examples to which we apply the above results: ultrafilter extensions of coalgebras [KKP05]; and a generalised determinisation procedure for probabilistic automata [BSS21; GP20; Goy21]. For the application of our results, some investigation of the definitions of the liftings involved is required. Theorem 3.3.11 then shows that a step common to both constructions (forgetting algebraic structure) does not affect our notion of bisimilarity. Proposition 3.3.13 recovers a result in the context of ultrafilter extensions showing that the (topological) closure of certain bisimulations are again bisimulations [BFV10].

Due to the generality of our approach and simplicity of the obtained conditions, our proofs of the above properties are simpler than those already present in the literature.

Outline Section 3.2 presents (invertible) steps, the relation to weak liftings and distributive laws, and a range of examples. In Section 2.4, we recall the standard notion of coalgebraic bisimilarity, defined via relation lifting. Section 3.3 contains the main results on preservation and reflection of bisimilarity. In Section 3.4 we discuss applications and instances of these results. We discuss other notions of bisimulation as well as future work in Section 3.5.

3.2 Forward and Backward Steps

We briefly present the required theory of steps, first termed as such in [RJL21]. This structure occurs already in work on coalgebraic modal logic [KKP04; BK05; PMW06; Kli07; CJ14; Lev15] where a step gives the one-step semantics of a logic. In existing work, only what we call a *forward* step is considered. Here, we also speak of *backward* steps, being arrows in the opposite direction.

In the sequel, we will usually assume that the backward step is a right inverse of the forward step, in which case the forward step will be referred to as an *invertible step*.

Next, we recall how such steps give rise to liftings of functors between categories of coalgebras and further, when the adjunction underlying the steps can also be lifted to coalgebras [HJ98]. Finally, we present examples of invertible steps from the literature, which we return to in later sections.

3.2.1 Invertible Steps

The basic setting of interest in this work consists of the following:

Definition 3.2.1: Given an adjunction $P \dashv Q\colon \mathcal{D} \to \mathcal{C}$ and endofunctors $B\colon \mathcal{C} \to \mathcal{C}$ and $L\colon \mathcal{D} \to \mathcal{D}$ as in the diagram

$$\begin{tikzcd} \mathcal{C} \arrow[loop left, "B"] \arrow[r, bend left, "P"] \arrow[r, phantom, "\perp"] & \mathcal{D} \arrow[l, bend left, "Q"] \arrow[loop right, "L"] \end{tikzcd}, \tag{3.1}$$

a *(forward) step* is a natural transformation $\delta\colon BQ \to QL$. A *backward step* is simply a natural transformation $\iota\colon QL \to BQ$ going the other way. If, moreover, $\delta \circ \iota = \mathrm{id}$ then we call δ an *invertible step* (with right inverse ι). Finally, if δ witnesses an isomorphism then we call it an *isomorphic step*. ◂

Notice the asymmetry in the definition of invertible step: ι is always assumed to be a *right* inverse of δ. These invertible steps are the main focus of this chapter. Before we treat two examples, note that (full) distributive laws as used in earlier work on determinisation for trace semantics [JSS15] yield examples of invertible steps. Such distributive laws are known to correspond to liftings of monads to Eilenberg-Moore categories. Such liftings are witnessed by equalities/natural isomorphisms $U\widetilde{\mathcal{S}} = \mathcal{S}U$ (see also Chapter 2). These are of course a special case of invertible steps.

Example: The Vietoris Monad as a weak lifting

Our first example of an invertible step arises as a weak lifting as defined by Garner. The notion of *weak* lifting replaces the equality in the lifting exactly by an invertible step $\delta\colon BQ \to QL$, while also assuming that the category $\mathcal{D}$ is some category of Eilenberg-Moore algebras of a monad on $\mathcal{C}$.

Garner shows that the Vietoris monad $\mathcal{V}$ on the category CHaus of compact Hausdorff spaces is an example of this. Namely, it is the weak lifting of the powerset monad on Set [Gar20] (see Chapter 2 for a more detailed introduction to weak liftings). This result relies on the fact that CHaus is equivalent to the category of algebras $\mathcal{EM}(\beta)$ for the ultrafilter monad β [Man69].

On objects, the Vietoris monad sends a compact Hausdorff space $\boldsymbol{X} = (X, U)$ to the *Vietoris hyperspace* [Vie22] $\mathcal{V}\boldsymbol{X}$ with underlying set $\mathcal{V}X := \{C \subseteq X \mid C \text{ closed in } \boldsymbol{X}\}$ and topology generated by the subbasis consisting of the following sets defined for each $C \in \mathcal{V}X$:

$$C^+ = \{A \in \mathcal{V}X \mid A \cap X = \varnothing\} \quad \text{and} \quad C^- = \{A \in \mathcal{V}X \mid A \not\subseteq X\} \tag{3.2}$$

On maps, $\mathcal{V}$ acts as the direct image (see also [Wyl06]). To place this into the setting of weak liftings, we consider the following:

$$\begin{tikzcd} \mathsf{Set} \arrow[loop left, "\mathcal{P}"] \arrow[r, bend left, "\mathcal{F}"] \arrow[r, phantom, "\perp"] & \mathcal{EM}(\beta) \arrow[l, bend left, "U"] \arrow[loop right, "\mathcal{V}"] \end{tikzcd} \tag{3.3}$$

where $\mathcal{F} \dashv U$ is the Eilenberg-Moore adjunction of β. The weak lifting provided by Garner then consists of natural transformations

$$U\mathcal{V} \xrightarrow{\iota} \mathcal{P}U \xrightarrow{\delta} U\mathcal{V} \tag{3.4}$$

satisfying $\delta \circ \iota = \mathrm{id}$ (and additional axioms discussed in Chapter 2, which are not required here). Notice that δ is an invertible step, with right inverse ι. As shown by Garner, a component $\delta_{\boldsymbol{X}} : \mathcal{P}U\boldsymbol{X} \to U\mathcal{V}\boldsymbol{X}$ sends each subset $S \in \mathcal{P}U\boldsymbol{X}$ to its topological closure. The components of ι simply include the closed subsets into the powerset.

Example: Convex Labelled Transition Systems

Another example of an invertible step occurs in a procedure for "determinising" probabilistic automata (PAs) into belief state transformers introduced in [BSS17; BSS21]. PAs combine non-determinism (in the sense of the powerset) with probabilities. The determinisation removes the probabilistic branching by constructing states consisting of distributions representing our "belief" about which state of the original PA we can be in. It was shown in [GP20] and the thesis of Goy [Goy21] that this is an instance of a more general determinisation procedure induced by a canonical weak distributive law. These correspond in general to a weak lifting (also shown by Garner [Gar20]), again being what we call an invertible step.

The presentation in [GP20] uses weak distributive laws of a monad $\mathcal{T}$ over a *monad* $\mathcal{S}$, while Goy in [Goy21] uses the simpler setting of a weak distributive law of a monad $\mathcal{T}$ over a *functor* S (see Chapter 2). We prefer the latter presentation here, as the corresponding weak lifting has no additional axioms; it is exactly an invertible step (see Definition 2.2.6).

We will consider the example of the Eilenberg-Moore adjunction of the distribution monad $\mathcal{D}$, and a weak lifting of the composite functor $\mathcal{R} \circ \mathcal{P}$, where we use $\mathcal{R}$ to denote the functor $(-)^{\Sigma}$ with Σ some set of labels. In a diagram:

$$\begin{tikzcd} \mathrm{Set} \arrow[loop left, "\mathcal{RP}"] \arrow[r, bend left, "\mathcal{F}"] \arrow[r, phantom, "\perp"] & \mathcal{EM}(\mathcal{D}) \arrow[l, bend left, "U"] \arrow[loop right, "\widetilde{\mathcal{RP}}"] \end{tikzcd} \tag{3.5}$$

Coalgebras for $\mathcal{RP}$ model standard labelled transition systems, while coalgebras for the weak lifting can again be thought of as labelled transition systems, but now with convex structure on the states. Such coalgebras with free $\mathcal{D}$-algebras as their state space arise after applying the first step of the generalised determinisation procedure for probabilistic automata introduced earlier, and are transformed back to LTSs by the second step. We wish to apply our conditions for preservation and reflection to this setting, especially the second step of the construction. This requires the existence of an invertible step.

We obtain this invertible step here by its correspondence to a weak distributive law $\lambda^{\mathcal{RP}} : \mathcal{DRP} \to \mathcal{RPD}$, which Goy shows can be obtained by combining weak distributive laws $\lambda^{\mathcal{P}} : \mathcal{DP} \to \mathcal{PD}$ and $\lambda^{\mathcal{R}} : \mathcal{DR} \to \mathcal{RD}$ in [Goy21, Ex 4.5]. We give the definitions of the latter two, before showing explicitly how they can be combined, and reach a concise description of the corresponding invertible step.

Definition 3.2.2 ([Goy21, Thm. 3.12]): A component $\lambda^{\mathcal{P}}_X : \mathcal{DP}X \to \mathcal{PD}X$ of the weak distributive law $\lambda^{\mathcal{P}}$, maps a distribution $\varphi : \mathcal{P}X \to [0,1]$ to

$$\left\{\mu_X\left(\sum_{U\subseteq X} \varphi(U)\cdot|\psi_U\rangle\right) \mid \forall U.\, \psi_U \in \mathcal{D}X \wedge \mathrm{supp}(\psi_U) \subseteq U\right\} \tag{3.6}$$ ◂

Definition 3.2.3: A component $\lambda^{\mathcal{R}}_X : \mathcal{DR}X \to \mathcal{RD}X$ of the weak distributive law $\lambda^{\mathcal{R}}$ maps a distribution $\varphi : X^{\Sigma} \to [0,1]$ to

$$\lambda a. \sum_{h\,:\,\Sigma\to X} \varphi(h)\cdot|h(a)\rangle \tag{3.7}$$ ◂

Proposition 3.2.4: The composition $\mathcal{R}\lambda^{\mathcal{P}} \circ \lambda^{\mathcal{R}}\mathcal{P}$ yields a natural transformation $\lambda^{\mathcal{RP}} : \mathcal{DRP} \to \mathcal{RPD}$. Concretely, a distribution $\varphi : (\mathcal{P}X)^{\Sigma} \to [0,1]$ is mapped to

$$\lambda a. \left\{\mu_X\left(\sum_{h\,:\,\Sigma\to\mathcal{P}X} \varphi(h)\cdot|\psi_h\rangle\right) \mid \forall h.\, \psi_h \in \mathcal{D}X \wedge \mathrm{supp}(\psi_h) \subseteq h(a)\right\} \tag{3.8}$$ ◂

The given concrete description of $\lambda^{\mathcal{RP}}$ is a slight reformulation of [Goy21, Eq. 4.49].

Using the constructions from [Gar20] (which apply also to the monad-functor setting), we can now find the invertible step we require to apply our later results to this setting.

Proposition 3.2.5: There is an invertible step $\delta : \mathcal{RP}U \to U\widetilde{\mathcal{RP}}$ with right inverse $\iota : U\widetilde{\mathcal{RP}} \to \mathcal{RP}U$.

A component $\delta_{\boldsymbol{A}}$, for $\boldsymbol{A} = (X, \alpha)$, sends a function $h : \Sigma \to \mathcal{P}X$ to

$$\lambda a.\, \mathrm{conv}_\alpha(h(a)) \tag{3.9}$$

with

$$\mathrm{conv}_\alpha(V \subseteq X) = \{\alpha(\varphi) \mid \varphi \in \mathcal{D}X \wedge \mathrm{supp}(\varphi) \subseteq V\} \tag{3.10}$$

Components $\iota_{\boldsymbol{A}}$ in turn act simply as the identity. ◂

Proof. As shown in [Gar20], the components of δ and ι arise via the following image-factorisation in Set:

$$\begin{tikzcd} \mathcal{RP}U(X,\alpha) \arrow[r, two heads, "\delta_A"] \arrow[rr, bend right, "\mathcal{RP}\alpha\circ\lambda_X\circ\eta_{\mathcal{RP}X}"'] & U\widetilde{\mathcal{RP}}(X,\alpha) \arrow[r, tail, "\iota_A"] & \mathcal{RP}U(X,\alpha) \end{tikzcd} \tag{3.11}$$

with $\boldsymbol{A} = (X, \alpha) \in \mathcal{EM}(\mathcal{D})$. It is also shown in [Gar20] that the splitting yields transformations such that $\delta \circ \iota = \mathrm{id}$.

By the definition of the image factorisation, $\delta_{\boldsymbol{A}}(h : \Sigma \to \mathcal{P}X)$ is equal to

$$\lambda a.\, \{\alpha(\psi_h) \mid \psi_h \in \mathcal{D}X \wedge \mathrm{supp}(\psi_h) \subseteq h(a)\} \tag{3.12}$$

so that the image of $h : \Sigma \to \mathcal{P}X$ under $\delta_{\boldsymbol{A}}$ is indeed as given above.

The components of ι act as the identity as they are image inclusions. ■

In [Gar20; Goy21] the transformations of the invertible step are obtained as an instance of a more general construction, namely: splitting the idempotent $\mathcal{RP}\alpha \circ \lambda_X \circ \eta_{\mathcal{RP}X}$ in the category of $\mathcal{D}$-*semialgebras*, and then applying the forgetful functor. As we work over Set, we can do the splitting there instead. The fact that the splitting of an idempotent $\mathcal{D}$-semialgebra morphism coincides with the splitting in Set is already implicitly used in [Gar20].

The action of this weak lifting can also be seen as a composition of the actions of the weak liftings of $\mathcal{P}$ and $\mathcal{R}$ induced by the earlier weak distributive laws. That of $\mathcal{P}$ performs the convex closure, while that of $\mathcal{R}$ applies an h under the algebra structure and abstraction over Σ. The construction of composite (or *iterated*) distributive laws is shown in [Goy21, Secs. 2.4 & 4.2]. However, it seems to be folklore that the weak lifting of (in this case) $\mathcal{RP}$ to $\widetilde{\mathcal{RP}}$ induced by a weak distributive law $\mathcal{DRP} \to \mathcal{RPD}$, is the composition of the weak liftings $\widetilde{\mathcal{R}}$ and $\widetilde{\mathcal{P}}$ induced by weak distributive laws $\mathcal{DR} \to \mathcal{RD}$ and $\mathcal{DP} \to \mathcal{PD}$. We prove this here for completeness.

Lemma 3.2.6: Let $\lambda^{\mathcal{R}} : \mathcal{DR} \to \mathcal{RD}$ and $\lambda^{\mathcal{P}} : \mathcal{DP} \to \mathcal{PD}$ be weak distributive laws inducing $\widetilde{\mathcal{R}}, \widetilde{\mathcal{P}} : \mathcal{EM}(\mathcal{D}) \to \mathcal{EM}(\mathcal{D})$. Then the composite weak distributive law $\lambda = \lambda^{\mathcal{R}}\mathcal{P} \circ \mathcal{R}\lambda^{\mathcal{P}}$ induces $\widetilde{\mathcal{RP}}$ such that $\widetilde{\mathcal{RP}} = \widetilde{\mathcal{R}} \circ \widetilde{\mathcal{P}}$. ◂

Proof. The weak lifting $\widetilde{\mathcal{RP}}((X, \alpha))$ of the composite functor $\mathcal{RP}$ is obtained by splitting the idempotent $\mathcal{RP}\alpha \circ \lambda_X \circ \eta_{\mathcal{RP}X}$.

The composite weak lifting $\widetilde{\mathcal{R}}\widetilde{\mathcal{P}}((X, \alpha))$ is instead obtained by splitting $\mathcal{R}\widetilde{\mathcal{P}}\alpha \circ \lambda^{\mathcal{R}}_{\mathcal{P}X} \circ \eta_{\mathcal{RP}X}$, where we write $\widetilde{\mathcal{P}}\alpha$ for the structure map of the algebra $\widetilde{\mathcal{P}}(X, \alpha)$.

We can thus show that the weak liftings coincide by showing that the following diagram commutes:

$$
\begin{tikzcd}
\mathcal{RP}X \arrow[r, "\eta_{\mathcal{RP}X}"] & \mathcal{DRP}X \arrow[r, "\lambda_X"] \arrow[dr, "\lambda^{\mathcal{R}}_{\mathcal{P}X}"'] & \mathcal{RPD}X \arrow[r, "\mathcal{RP}\alpha"] & \mathcal{RP}X \\
& & \mathcal{RDP}X \arrow[u, "\mathcal{R}\lambda^{\mathcal{P}}_X"'] \arrow[ur, "\mathcal{R}\widetilde{\mathcal{P}}\alpha"'] &
\end{tikzcd}
\tag{3.13}
$$

The left triangle is the definition of λ. The right triangle is $\mathcal{R}$ applied to the definition of $\widetilde{\mathcal{P}}\alpha$. ■

Remark 3.2.7: Later, we will also refer to the *convex powerset functor* $\mathcal{P}_c$. This is exactly the weak lifting of $\mathcal{P}$ corresponding to $\lambda^{\mathcal{P}}$, producing an algebra on the set of convex subsets of a convex algebra (a set being convex if it is closed under convex combinations). This functor is found in [GP20; BSS21] and [Goy21, Chapter 3], which deal with the combination of distributions with non-determinism in more detail. Note that in [BSS21], the empty set is excluded from the *non-empty* powerset functor $\mathcal{P}_{ne}$ and convex powerset functor $\mathcal{P}_c$. The fact that there is a unique *functorial* way to extend convex algebras with a point, shown in [SW17], then allows the extension to $\mathcal{P}_{ne} + 1$ and $\mathcal{P}_c + 1$. ◂

3.2.2 Step-induced liftings

In the setting of Definition 3.2.1, there is a bijective correspondence between a step and its *mate* $\hat{\delta} : PB \to LP$ given by

$$
PB \xrightarrow{PB\eta} PBQP \xrightarrow{P\delta P} PQLP \xrightarrow{\varepsilon LP} LP
\tag{3.14}
$$

with η and ε the unit and counit of the adjunction $P \dashv Q$ (see also [Lei04; KS74]). The following lemma will turn out to be useful in relating a step and its mate:

Lemma 3.2.8: For a forward step $\delta\colon BQ \to QL$, we have $\varepsilon L \circ P\delta = L\varepsilon \circ \hat{\delta}Q$ and $\delta P \circ B\eta = Q\hat{\delta} \circ \eta B$. ◂

Proof. See, e.g., [RJL21]. ∎

Now, the mate and the backward step allow us to define liftings of P and Q to functors between the categories of coalgebras for B and L.

Definition 3.2.9: Given steps $\delta\colon BQ \to QL$ and $\iota\colon QL \to BQ$, the *step-induced coalgebra liftings* $\overline{P}\colon \mathrm{Coalg}(B) \to \mathrm{Coalg}(L)$ and $\overline{Q}\colon \mathrm{Coalg}(L) \to \mathrm{Coalg}(B)$ of P and Q are defined by

$$\gamma\colon X \to BX \quad \mapsto \quad \hat{\delta}_X \circ P\gamma\colon PX \to LPX \tag{3.15}$$

$$\xi\colon Y \to LY \quad \mapsto \quad \iota_Y \circ Q\xi\colon QY \to BQY \tag{3.16}$$

on objects and act as P and Q on arrows. This is well-defined due to functoriality of P and Q and naturality of $\hat{\delta}$ and ι. ◂

It is shown in [HJ98, Theorem 2.14] that when δ and ι form an isomorphism, the adjunction $P \dashv Q$ lifts to an adjunction $\overline{P} \dashv \overline{Q}$ between the step-induced liftings. For our purposes it will be useful to split the isomorphism condition into the cases where ι is a left or right inverse of δ, mostly for the right inverse case as this is exactly an invertible step.

Lemma 3.2.10: If $\delta \circ \iota = \mathrm{id}$, then the counit $\varepsilon\colon PQ \to \mathrm{Id}$ of the adjunction $P \dashv Q$ lifts to a natural transformation $\overline{\varepsilon}\colon \overline{P}\,\overline{Q} \to \mathrm{Id}$. If $\iota \circ \delta = \mathrm{id}$, then the unit $\eta\colon \mathrm{Id} \to QP$ of the adjunction lifts to a natural transformation $\overline{\eta}\colon \mathrm{Id} \to \overline{Q}\,\overline{P}$. ◂

Proof. The first part is shown in Fig. 3.1. The second part is shown in Fig. 3.2. The parts of the diagrams marked *Lem.* commute by Lemma 3.2.8. ∎

The combination of these two liftings gives us the lifting of the adjunction.

Corollary 3.2.11 ([HJ98, Thm. 2.14]): If δ and ι form an isomorphism, then $\overline{P} \dashv \overline{Q}$. ◂

Proof. We have $\delta \circ \iota = \mathrm{id}$ and $\iota \circ \delta = \mathrm{id}$ due to the isomorphism, so both the unit and counit lift. The required identities for the components of the liftings hold by the adjunction $P \dashv Q$. ∎

$$\begin{tikzcd}
PQY \arrow[rr, "\varepsilon_Y"] \arrow[d, "PQ\xi"'] & \text{nat. } \varepsilon & Y \arrow[ddd, "\xi"] \\
PQLY \arrow[rrdd, bend left, "\varepsilon_{LY}"] \arrow[d, "P\iota_Y"'] & \delta \circ \iota = \mathrm{id} & \\
PBQY \arrow[r, "P\delta_Y"] \arrow[d, "\hat{\delta}_{QY}"'] & PQLY \arrow[rd, "\varepsilon_{LY}"'] & \\
LPQY \arrow[rr, "L\varepsilon_Y"'] & \text{Lem.} & LY
\end{tikzcd}$$

Figure 3.1: Lifting ε

$$\begin{tikzcd}
X \arrow[rr, "\eta_X"] \arrow[ddd, "\gamma"'] & \text{nat. } \eta & QPX \arrow[d, "QP\gamma"] \\
& & QPBX \arrow[d, "Q\hat{\delta}_X"] \\
& BQPX \arrow[r, "\delta_{PX}"] & QLPX \arrow[d, "\iota_{PX}"] \\
BX \arrow[rruu, bend left, "\eta_{BX}"] \arrow[ru, "B\eta_X"'] \arrow[rr, "B\eta_X"'] & \iota \circ \delta = \mathrm{id} & BQPX
\end{tikzcd}$$

Figure 3.2: Lifting η

In such a situation, $\overline{Q}$ (being a right adjoint) preserves the final coalgebra for L (the limit of the empty diagram in $\mathrm{Coalg}(L)$) when this exists. However, there are a number of known examples where the step is not an isomorphism; instead we only have a one-sided inverse. We consider, in particular, these invertible steps and now continue with examples of this setting.

We return to the two examples of invertible steps, and discuss the liftings to coalgebras they induce. We will later apply our results on preservation and reflection of bisimilarity/behavioural equivalence to these examples.

Example: The Vietoris Monad continued

Returning to the example of the Vietoris monad, it turns out that the invertible step in that setting gives rise to ultrafilter extensions of Kripke frames. In modal logic, ultrafilter extensions [BRV01; Gol74; Ben79] are a construction taking a Kripke frame with state space W and forming a new Kripke frame with states being ultrafilters over W. The central motivation for this is in "bisimilarity-somewhere-else" results: two states are modally equivalent iff they are bisimilar in the ultrafilter extension.

The transition relation R of a Kripke frame (W, R) is equivalently a powerset coalgebra $\gamma\colon X \to \mathcal{P}X$ with $\gamma(x) = \{x' \mid x\ R\ x'\}$. Now, the composition

$$\mathrm{Coalg}(\mathcal{P}) \xrightarrow{\overline{\mathcal{F}}} \mathrm{Coalg}(\mathcal{V}) \xrightarrow{\overline{U}} \mathrm{Coalg}(\mathcal{P}) \tag{3.17}$$

of the step-induced coalgebra liftings precisely yields the ultrafilter extension of a Kripke frame. The first stage $\overline{\mathcal{F}}$ is the actual extension, which turns the Kripke frame into a $\mathcal{V}$-coalgebra whose state space is the free β-algebra on X. The second stage $\overline{U}$ turns this back into a Kripke frame, i.e., a powerset coalgebra in Set. Applying these two stages to a coalgebra $\gamma\colon X \to \mathcal{P}X$, we obtain a coalgebra with state space $U \circ \mathcal{F}(X) = \beta X$, the set of ultrafilters over X. The transition function, which we call $\overline{\gamma}$ here, is the composition

$$\beta X \cong \mathcal{U}\mathcal{F}X \xrightarrow{\mathcal{U}\mathcal{F}\gamma} \mathcal{U}\mathcal{F}\mathcal{P}X \xrightarrow{\mathcal{U}\hat{\delta}_X} \mathcal{U}\mathcal{V}\mathcal{F}X \xrightarrow{\iota_{\mathcal{F}X}} \mathcal{P}\mathcal{U}\mathcal{F}X \cong \mathcal{P}\beta X \tag{3.18}$$

It is shown in [KKP05] that this coincides with the classical definition of ultrafilter extension whereby $v \in \overline{\gamma}(u)$ for $u, v \in \beta X$ iff the set

$$m(Y) := \{x \in X \mid \exists y \in Y.\, y \in \gamma(x)\} \tag{3.19}$$

is contained in u for all $Y \in v$.

Their approach uses a transformation $\pi\colon \mathcal{F}\mathcal{P}\mathcal{U} \to \mathcal{V}$ to obtain the above coalgebra as the composition $\pi_{\mathcal{F}X} \circ \mathcal{F}\mathcal{P}\eta_X \circ \mathcal{F}\gamma$. This can be seen to match the above definition via the bijective correspondence between transformations $\mathcal{F}\mathcal{P}\mathcal{U} \to \mathcal{V}$ and $\mathcal{F}\mathcal{P} \to \mathcal{V}\mathcal{F}$, as given in [RJL21].

Remark 3.2.12: In [KKP05], ultrafilter extensions are developed more generally for coalgebras for a functor $B\colon \mathsf{Set} \to \mathsf{Set}$, via the duality between Boolean algebras and Stone spaces. In fact, since both $\mathcal{V}$ and the left adjoint $\mathcal{F}$ restrict to the category Stone of Stone spaces, the invertible step δ, ι restricts to an invertible step in the restriction of the above adjunction to Stone. When we return to this example in Section 3.4, we will thus replace the category $\mathcal{EM}(\beta)$ in Eq. (3.3) with Stone. ◂

Example: Probabilistic Automata

The invertible step corresponding to the weak distributive law $\lambda^{\mathcal{R}\mathcal{P}}$ presented in Proposition 3.2.5, can be applied in a generalised determinisation procedure studied for probabilistic automata in [GP20; Goy21]. This is adapted from a general construction in the setting of an Eilenberg-Moore adjunction $\mathcal{F} \dashv U\colon \mathcal{EM}(\mathcal{T}) \to \mathcal{C}$, and monad $B\colon \mathcal{C} \to \mathcal{C}$, which further assumes a (full) lifting of B to algebras (corresponding to a distributive law of $\mathcal{T}$ over B). It was observed in [JSS15] that the determinisation then arises as the following two-step process (cf. Eq. (3.17)):

$$\mathsf{Coalg}_{\mathcal{C}}(B\mathcal{T}) \xrightarrow{\overline{\mathcal{F}}} \mathsf{Coalg}_{\mathcal{EM}(\mathcal{T})}(\overline{B}) \xrightarrow{\overline{U}} \mathsf{Coalg}_{\mathcal{C}}(B) \tag{3.20}$$

The approach based on *weak* distributive laws of Goy and Petrisan instead uses an invertible step (weak lifting) $\delta\colon BU \to U\widetilde{B}$. This gives rise to a very similar two-step diagram, in which $\widetilde{B}$ is now a weak lifting:

$$\mathsf{Coalg}_{\mathcal{C}}(B\mathcal{T}) \xrightarrow{\overline{\mathcal{F}}} \mathsf{Coalg}_{\mathcal{EM}(\mathcal{T})}(\widetilde{B}) \xrightarrow{\overline{U}} \mathsf{Coalg}_{\mathcal{C}}(B) \tag{3.21}$$

The second functor $\overline{U}$ is simply the lifting of U induced by the right-inverse of δ. The first ($\overline{\mathcal{F}}$) is a variation of a step-induced lifting (notice that it takes BT-coalgebras rather than B-coalgebras as input), mapping a coalgebra $\gamma\colon X \to BTX$ to

$$\mathcal{F}X \xrightarrow{\mathcal{F}\gamma} \mathcal{F}BU\mathcal{F}X \xrightarrow{\hat{\delta}_{U\mathcal{F}X}} \widetilde{B}\mathcal{F}U\mathcal{F}X \xrightarrow{\widetilde{B}\varepsilon_{\mathcal{F}X}} \widetilde{B}\mathcal{F}X \tag{3.22}$$

where ε is the counit of the Eilenberg-Moore adjunction. In fact, this can be viewed as a step-induced lifting for BT which arises by composing δ and the counit, see [RJL21].

As explained in [GP20, Sec. 5] and [Goy21, Sec. 4.3], the weak lifting which corresponds to the invertible step $\delta\colon \mathcal{R}\mathcal{P}U \to U\widetilde{\mathcal{R}\mathcal{P}}$ elaborated in Proposition 3.2.5 induces the lifting $\overline{\mathcal{F}}$ of (3.21). This lifting indeed gives the transformation of a probabilistic automaton into a belief-state transformer in the category $\mathcal{EM}(\mathcal{D})$. The second step is then to transfer the obtained belief state transformer back to Set with the lifting of U induced by the right inverse $\iota\colon U\mathcal{R}\mathcal{P} \to U\widetilde{\mathcal{R}\mathcal{P}}$ of δ. As shown in [BSS21] for the setting of a *lax* lifting and later recovered from our abstract theory for invertible steps (Section 3.4), this second step yields a system with the same behaviour.

In Section 3.4 we will return to the two examples just presented and show how we can apply the techniques developed in the coming section (3.3) to obtain preservation and reflection of bisimilarity by the step-induced liftings of the involved forgetful functors. To further show the applicability of our abstract framework, we also recover a number of other results from works on bisimulations for coalgebras of the Vietoris functor [BFV10] and the convex powerset functor [BSS21].

3.3 Preserving and Reflecting Bisimilarity

In this section we show that, in the presence of an invertible step, bisimilarity is preserved and reflected by the step-induced lifting of the right adjoint, given some further mild conditions. Aside from this main motivating result for right adjoints, these conditions will combine with other properties shown in Lemma 3.3.2 to recover a number of existing results relating various notions of equivalence on either side of an adjunction (Section 3.4).

Our approach is as follows:

- We show how the setting of Definition 3.2.1 (an adjunction with endofunctors on each side) lifts to categories of relations.

- We show that the steps themselves also lift to these categories.

- We then restrict to lattices of relations on fixed objects, allowing the application of Lemma 3.3.2 to obtain our conditions for preservation and reflection by the step-induced liftings of adjoint functors.

Throughout this section we assume the setting of Definition 3.2.1 with the additional properties that $\mathcal{C}$ and $\mathcal{D}$ are complete, well-powered categories with logical factorisation systems $(\mathcal{E}_1, \mathcal{M}_1), (\mathcal{E}_2, \mathcal{M}_2)$ respectively such that $P(\mathcal{E}_1) \subseteq \mathcal{E}_2$, and subobjects defined with respect to the set of *abstract* monos. We further assume an invertible step $\delta \colon BQ \to QL$ with right inverse $\iota \colon QL \to BQ$.

We first fix a definition of preservation and reflection of post- and greatest-fixed points in lattices. Lemma 3.3.2 then summarises the sufficient conditions of monotone maps we use to speak of preservation and reflection of bisimulations and bisimilarity and later (in Chapter 4) of coinductive invariants and predicates.

3.3.1 Preservation and reflection of fixpoints

As discussed in Section 2.4, bisimulations can be defined as post-fixed points of a monotone map $f \colon \Gamma \to \Gamma$ on a poset Γ consisting of relations on the carrier X of some coalgebra. Similarly, bisimulations on a coalgebra to which a step-induced lifting (of Q say) has been applied are post-fixed points of a monotone map $g \colon \Delta \to \Delta$ on the poset of relations on QX. The general picture is:

$$\begin{tikzcd} \Delta \arrow[loop left, "b"] \arrow[r, bend left, "f"] \arrow[r, phantom, "\perp"] & \Gamma \arrow[l, bend left, "g"] \arrow[loop right, "l"] \end{tikzcd} \tag{3.23}$$

We can thus speak of preservation and reflection of bisimulations and bisimilarity by speaking of preservation and reflection of fixpoints. In the following definition, we fix some necessary terminology.

Definition 3.3.1: Let Γ and Δ be posets, and $l \colon \Gamma \to \Gamma$, $b \colon \Delta \to \Delta$ monotone maps. A monotone map $g \colon \Gamma \to \Delta$ *preserves post-fixed points* if $x \le l(x)$ implies $g(x) \le b(g(x))$. It *reflects post-fixed points* if the converse implication holds. Furthermore, we will speak of g *preserving the greatest fixed point* if $g(\nu l) = \nu b$. ◀

Now, we summarise results concerning conditions for such preservation and reflection properties. Those relating greatest fixed points can be found already in, e.g., [DP02]. The other items are essentially instances of step-induced liftings [RJL21].

Lemma 3.3.2: Let Δ and Γ be lattices, and f, g, b, l monotone maps as in Eq. (3.23). Then:

1. if $gl \leq bg$, then
   a) $g(\nu l) \leq \nu b$;
   b) for any $p \in \Gamma$ we have $p \leq l(p) \implies g(p) \leq bg(p)$;
2. if $bg \leq gl$, then
   a) $g(\nu l) \geq \nu b$;
   b) if g is order-reflecting for any $p \in \Gamma$ we have
   $$p \leq l(p) \impliedby g(p) \leq bg(p);$$
   c) $\nu l \geq f(\nu b)$;
   d) for any $q \in \Delta$ we have $q \leq b(q) \implies f(q) \leq lf(q)$. ◂

Proof. For Item 1a, we have $g(\nu l) = gl(\nu l) \leq bg(\nu l)$ so that $g(\nu l)$ is a post-fixed point of b and thus below νb.

For Item 1b, we have $p \leq l(p) \implies g(p) \leq gl(p) \leq bg(p)$ with the last inequality holding by assumption.

For Item 2a, the consequent is equivalent to $\nu l \geq f(\nu b)$ by the adjunction. It is therefore sufficient to prove that $f(\nu b)$ is a post-fixed point of l. Indeed, we have

$$\begin{array}{lrl}
\text{(unit of adjunction)} & & \nu b \leq gf(\nu b) \\
(b \text{ monotone}) & \implies & b(\nu b) \leq bgf(\nu b) \\
\text{(def. fixpoint)} & \implies & \nu b \leq bgf(\nu b) \\
\text{(assum.)} & \implies & \nu b \leq glf(\nu b) \\
\text{(adjunction)} & \implies & f(\nu b) \leq lf(\nu b)
\end{array}$$

For Item 2b, we have $g(p) \leq bg(p) \leq gl(p) \implies p \leq l(p)$ where the implication is exactly order-reflection.

For Item 2c, we have already noted that this is equivalent to the inequality of Item 2a, by the adjunction.

Finally, for Item 2d, we have $q \leq b(q) \implies f(q) \leq fb(q) \leq lf(q)$. The final inequality holds due to the chain $fb \leq fbgf \leq fglf \leq lf$. ■

Now, if $bg = gl$ in the setting of (3.23), the inequalities combine to give $g(\nu l) = \nu b$. This has been shown more generally in the context of coalgebras in [HJ98], where the equality $bg = gl$ is instead a natural isomorphism

$BG \xrightarrow{\sim} GL$, with F, G, B, L functors. It is shown in *op. cit.* that this allows the lifting of the adjunction $F \dashv G$ to coalgebras (generalising post-fixed points) so that the right adjoint preserves the final coalgebra (generalising the greatest fixed point) as right adjoints preserve limits.

3.3.2 From adjunctions to lattices

In the setting of Eq. (3.1), and given a coalgebra $\xi : X \to LX$, the lattices and monotone maps which can be used to speak of preservation and reflection by the right adjoint are those of the following diagram (containing the as-yet-unknown map f):

$$\begin{tikzcd} \mathsf{Rel}_{QX} \arrow[loop left, "(\iota_X \circ Q\xi)^* \circ \mathsf{Rel}(B)_{QX}"] \arrow[r, bend left, "f"] \arrow[r, phantom, "\perp"] & \mathsf{Rel}_X \arrow[l, bend left, "\mathsf{Rel}(Q)_X"] \arrow[loop right, "\xi^* \circ \mathsf{Rel}(L)_X"] \end{tikzcd} \tag{3.24}$$

On the left, $\iota \circ Q\xi : QX \to BQX$ is the coalgebra obtained by applying the lifting of Q induced by ι to the coalgebra ξ. This setting thus allows us to compare bisimilarity on ξ and $\overline{Q}(\xi)$ via the relation lifting $\mathsf{Rel}(Q)$.

The missing ingredients to allow the application of Lemma 3.3.2 are thus:

1. A left adjoint for $\mathsf{Rel}(Q)_X$ (Lemma 3.3.7).

2. The equality $(\iota_X \circ Q\xi)^* \circ \mathsf{Rel}(B)_{QX} \circ \mathsf{Rel}(Q)_X = \mathsf{Rel}(Q)_X \circ \xi^* \circ \mathsf{Rel}(L)_X$, an instance of the combined inequalities $bg \leq gl$ and $gl \leq bg$ (Theorem 3.3.10).

3. Order-reflection of $\mathsf{Rel}(Q)_X$ (assumption; discussed at the end of this section).

To obtain the required adjunction between the fibres Rel_X and Rel_{QX}, we first establish the adjunction $\mathsf{Rel}(P) \dashv \mathsf{Rel}(Q)$ between the total relation categories. Given the above assumptions on the categories and their factorisation systems, we can lift the unit and counit of the adjunction $P \dashv Q$, using the transformations constructed in the following lemma.

Lemma 3.3.3: Let $F : \mathcal{C} \to \mathcal{D}$ and $G : \mathcal{D} \to \mathcal{E}$ be functors, with

$$\mathsf{Rel}(\mathcal{C}) \xrightarrow{\mathsf{Rel}(F)} \mathsf{Rel}(\mathcal{D}) \xrightarrow{\mathsf{Rel}(G)} \mathsf{Rel}(\mathcal{E}) \tag{3.25}$$

the corresponding relation liftings. Then the natural transformation (1) in Eq. (3.26) below exists. Furthermore, if G preserves abstract epis, then the natural transformation (2) exists.

(3.26)
$$\begin{tikzcd}
\mathsf{Rel}(\mathcal{C}) \arrow[rr, bend left, "\mathsf{Rel}(GF)"] \arrow[rr, bend right, "\mathsf{Rel}(G)\mathsf{Rel}(F)"'] \arrow[rr, phantom, "(1)\Downarrow\ \Uparrow(2)"] & & \mathsf{Rel}(\mathcal{E})
\end{tikzcd}$$

Further (when they exist) the transformations are above the identity. ◂

Proof. Consider:

(3.27)
$$\begin{tikzcd}
& \mathsf{Rel}(GF)(R) \arrow[rd, tail, "m_{GF}"] \arrow[d, phantom, "(1)"] & \\
GFR \arrow[ru, "e_{GF}"] \arrow[r, "GFr"] \arrow[d, "Ge_F"'] \arrow[rd, phantom, "(2)"] & GF(X \times X) \arrow[r, "\langle GF\pi_1, GF\pi_2 \rangle"] \arrow[d, "G\langle F\pi_1, F\pi_2 \rangle"] \arrow[rd, phantom, "(3)"] & GFX \times GFX \\
G(\mathsf{Rel}(F)(R)) \arrow[r, "Gm_F"] \arrow[rr, bend right, "e_G"'] & G(FX \times FX) \arrow[ru, "\langle G\pi_1, G\pi_2 \rangle"] \arrow[r, phantom, "(4)"] & \mathsf{Rel}(G)(\mathsf{Rel}(F)(R)) \arrow[u, tail, "m_G"']
\end{tikzcd}$$

We see that (1) and (4) commute by definition of relation lifting, (2) commutes as it is G applied to the square defining $\mathsf{Rel}(F)$, and (3) commutes by the properties of pairing. Then we obtain a unique map by diagonal fill-in as in:

(3.28)
$$\begin{tikzcd}
GFR \arrow[r, "e_{GF}"] \arrow[d, "Ge_F"] & \mathsf{Rel}(GF)(R) \arrow[dd, tail, "m_{GF}"] \arrow[ddl, dashed] \\
G(\mathsf{Rel}(F)(R)) \arrow[d, "e_G"] & \\
\mathsf{Rel}(G)(\mathsf{Rel}(F)(R)) \arrow[r, tail, "m_G"] & GFX \times GFX
\end{tikzcd}$$

The maps obtained in this way are above the identity because the lower right triangle commutes and we can add the identity on GFX to obtain the defining diagram of a map in Rel. Thus, maps constructed in this way are components of a natural transformation above the identity.

For the other direction, the components are obtained as diagonal fill-ins as in:

$$
\begin{tikzcd}
GFR \arrow[r, "Ge_F"] \arrow[d, "e_{GF}"] & G(\mathsf{Rel}(F)(R)) \arrow[r, "e_G"] & \mathsf{Rel}(G)(\mathsf{Rel}(F)(R)) \arrow[d, "m_G"] \arrow[dll, dashed] \\
\mathsf{Rel}(GF)(R) \arrow[rr, "m_{GF}"] & & GFX \times GFX
\end{tikzcd}
\tag{3.29}
$$

where Ge_F is abstract epi by assumption. ■

We note that the first part is in [Jac16, Exercise 4.4.6] and the result is proved for Set endofunctors in [Bon+17, Lemma 14.1]. The above lemma allows the lifting of the adjunction, which we note may also be obtainable from results on fibred adjunctions in [Jac01; Her92], but a direct proof is quite straightforward; the main idea is to use Lemma 3.3.3 together with preservation of abstract epis by P.

Lemma 3.3.4: The adjunction $P \dashv Q\colon \mathcal{D} \to \mathcal{C}$ lifts to relations, i.e., the following diagram is commutative:

$$
\begin{tikzcd}
\mathsf{Rel}(\mathcal{C}) \arrow[r, bend left, "\mathsf{Rel}(P)"] \arrow[r, phantom, "\perp"] \arrow[d, "p"'] & \mathsf{Rel}(\mathcal{D}) \arrow[l, bend left, "\mathsf{Rel}(Q)"] \arrow[d, "q"] \\
\mathcal{C} \arrow[r, bend left, "P"] \arrow[r, phantom, "\perp"] & \mathcal{D} \arrow[l, bend left, "Q"]
\end{tikzcd}
\tag{3.30}
$$

Further, the unit and counit of $\mathsf{Rel}(P) \dashv \mathsf{Rel}(Q)$ are above those of $P \dashv Q$. ◂

Proof. We define the unit and counit of the adjunction, and show that these satisfy the required conditions.

The unit $\overline{\eta}\colon \mathsf{Id} \to \mathsf{Rel}(Q)\mathsf{Rel}(P)$ is defined by

$$
\mathsf{Id}_{\mathsf{Rel}(\mathcal{C})} = \mathsf{Rel}(\mathsf{Id}_{\mathcal{C}}) \to \mathsf{Rel}(QP) \to \mathsf{Rel}(Q)\mathsf{Rel}(P) \tag{3.31}
$$

The counit $\overline{\varepsilon}\colon \mathsf{Rel}(P)\mathsf{Rel}(Q) \to \mathsf{Id}$ is defined by

$$
\mathsf{Rel}(P)\mathsf{Rel}(Q) \to \mathsf{Rel}(PQ) \to \mathsf{Rel}(\mathsf{Id}_{\mathcal{D}}) = \mathsf{Id}_{\mathsf{Rel}(\mathcal{D})} \tag{3.32}
$$

In each case we use 2-functoriality of Rel to lift the natural transformations (2-cells) η and ε to relations, together with Lemma 3.3.3 to construct the required component transformations. Note that P (being a left adjoint) preserves abstract epis by assumption.

Now, one of the required identities is that for any $R \rightarrowtail X \times X$, we have $\overline{\varepsilon}_{\mathsf{Rel}(F)(R)} \circ \mathsf{Rel}(F)(\overline{\eta}_R) = \mathsf{id}_{\mathsf{Rel}(F)(R)}$. By its definition, the left-hand map is above $\varepsilon_{FX} \circ F\eta_X = \mathsf{id}_{FX}$. Thus, it is a map in the fibre above FX which is a poset, so that arrows are unique.

The other identity involving $\mathsf{Rel}(G)$ holds similarly. ∎

The relation lifting defined in Section 2.4 allows us to define endofunctors $\mathsf{Rel}(B)$, $\mathsf{Rel}(L)$ in the context of the above adjunction:

$$
\begin{tikzcd}
\mathsf{Rel}(\mathcal{C}) \arrow[loop left, "\mathsf{Rel}(B)"] \arrow[r, bend left, "\mathsf{Rel}(P)"] \arrow[r, phantom, "\perp"] & \mathsf{Rel}(\mathcal{D}) \arrow[l, bend left, "\mathsf{Rel}(Q)"] \arrow[loop right, "\mathsf{Rel}(L)"]
\end{tikzcd}
\tag{3.33}
$$

In this setting, we may try to lift the steps δ or ι to this adjunction. It turns out that δ always lifts. For ι, there is a sufficient condition which is independent of ι itself: that Q preserves abstract epis. In both cases, this result follows essentially from Lemma 3.3.3.

Proposition 3.3.5: For a forward step δ and backward step ι, we have:

1. δ lifts to relations, i.e., there exists a natural transformation $\overline{\delta} \colon \mathsf{Rel}(B) \circ \mathsf{Rel}(Q) \to \mathsf{Rel}(Q) \circ \mathsf{Rel}(L)$ above δ.

2. If Q preserves abstract epis, then ι lifts to relations, i.e., there exists a natural transformation $\overline{\iota} \colon \mathsf{Rel}(Q) \circ \mathsf{Rel}(L) \to \mathsf{Rel}(B) \circ \mathsf{Rel}(Q)$ above ι. ◀

Proof. 1. By our assumption that left adjoints preserve abstract epis, we can lift the mate $\hat{\delta}$ (see Eq. (3.14)) to $\overline{\hat{\delta}} \colon \mathsf{Rel}(P) \circ \mathsf{Rel}(B) \to \mathsf{Rel}(L) \circ \mathsf{Rel}(P)$ using Lemma 3.3.3. Then, from the adjunction $\mathsf{Rel}(P) \dashv \mathsf{Rel}(Q)$, we obtain the lifting of δ.

2. The assumption that Q preserves abstract epis again allows us to apply Lemma 3.3.3 to lift ι. ∎

The condition that Q preserves abstract epis holds, e.g., in case it is the forgetful functor in an adjunction monadic over Set and we take the factorisation system $(\mathsf{RegEpi}, \mathsf{Mono})$ in each category, as is common when working with relations. Such factorisation systems exist in all regular categories [Bor94b, Thm. 2.1.3]. For example, Set and categories monadic over it. This, and the preservation of regular epis by $U \colon \mathcal{EM}(T) \to \mathrm{Set}$ are shown in [Bor94b, Thm. 4.3.5].

Preservation of abstract epis also holds for the forgetful functor in the Stone-Set adjunction as Stone is a (full) subcategory of CHaus, i.e., a regular epi in Stone maps to a regular epi in CHaus by the inclusion, and on to a regular epi in Set by the above.

The lifted steps $\overline{\delta}$ and $\overline{\iota}$ from Proposition 3.3.5 give step-induced liftings of $\mathsf{Rel}(P)$ and $\mathsf{Rel}(Q)$ between $\mathsf{Coalg}(\mathsf{Rel}(B))$ and $\mathsf{Coalg}(\mathsf{Rel}(L))$. While we primarily use the definition of coalgebraic bisimulations R on a coalgebra $\gamma\colon X \to BX$ as post-fixed points of $\gamma^* \circ \mathsf{Rel}(B)$, this is equivalent to R having a coalgebra structure for $\mathsf{Rel}(B)$ whose structure map is above γ (see, e.g., [Jac16, Def 6.2.1]). This means that the above lifted steps could already be used to capture preservation of bisimulations. However, it is less obvious what reflection means in this context and how to prove it. To have a coherent story of both preservation and reflection of bisimulations and preservation of bisimilarity by $\mathsf{Rel}(Q)$, we turn our attention to the fibres, as described in the beginning of this section.

As a consequence of Proposition 3.3.5 and of $\delta \circ \iota = \mathrm{id}$, we obtain the following result, which will later be used in the construction of a step on the fibres.

Theorem 3.3.6: Let δ be an invertible step with right inverse ι, and suppose Q preserves abstract epis. Then $\mathsf{Rel}(Q)_{LX} \circ \mathsf{Rel}(L)_X = \iota_X^* \circ \mathsf{Rel}(B)_{QX} \circ \mathsf{Rel}(Q)_X$. ◀

Proof. Both $\overline{\iota}$ and $\overline{\delta}$ exist by Proposition 3.3.5.

From the lifting $\overline{\iota}\colon \mathsf{Rel}(Q) \circ \mathsf{Rel}(L) \to \mathsf{Rel}(B) \circ \mathsf{Rel}(Q)$, we immediately obtain $\mathsf{Rel}(Q)_{LX} \circ \mathsf{Rel}(L)_X \leq \iota_X^* \circ \mathsf{Rel}(B)_{QX} \circ \mathsf{Rel}(Q)_X$ by universality of pullbacks.

For the other direction, we have $\mathsf{Rel}(B)_{QX} \circ \mathsf{Rel}(Q)_X \leq \delta_X^* \circ \mathsf{Rel}(Q)_{LX} \circ \mathsf{Rel}(L)_X$ by universality of pullbacks. Applying ι^* to both sides gives $\iota_X^* \circ \mathsf{Rel}(B)_{QX} \circ \mathsf{Rel}(Q)_X \leq \iota_X^* \circ \delta_X^* \circ \mathsf{Rel}(Q)_{LX} \circ \mathsf{Rel}(L)_X = \mathsf{Rel}(Q)_{LX} \circ \mathsf{Rel}(L)_X$, where the equality holds as $\delta \circ \iota = \mathrm{id}$ and $\iota_X^* \circ \delta_X^* = (\delta_X \circ \iota_X)^*$. ■

Adjoining the fibres

Now that we have an adjunction between the total categories, we use this to construct an adjunction between the fibres Rel_X and Rel_{QX} in order to satisfy the first condition at the beginning of this section. The usual restriction $\mathsf{Rel}(Q)_X$ of $\mathsf{Rel}(Q)$ to the fibre Rel_X will be the right adjoint, similarly to the adjunction obtained earlier. To map back into the fibre Rel_X, we post-compose $\mathsf{Rel}(P)_{QX}$ with $\coprod_\varepsilon$, the direct image functor obtained from the counit of the adjunction $\mathsf{Rel}(P) \dashv \mathsf{Rel}(Q)$. We note the similarity with results on fibred adjunctions in [Jac01], where only adjunctions over a single base category are considered.

Lemma 3.3.7: We have an adjunction $\coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX} \dashv \mathsf{Rel}(Q)_X : \mathsf{Rel}_X \to \mathsf{Rel}_{QX}$. ◂

Proof. We define the unit of the adjunction $R \to \mathsf{Rel}(Q)_X \circ \coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX}(R)$ as the transpose of the composition

$$\mathsf{Rel}(P)_{QX}(R) \xrightarrow{\eta} \varepsilon^*(\coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX}(R)) \xrightarrow{\bar{\varepsilon}} \coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX}(R) \tag{3.34}$$

We define the counit $\coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX} \circ \mathsf{Rel}(Q)_X(R) \to R$ via a number of correspondences. First

$$\coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX} \circ \mathsf{Rel}(Q)_X(R) \to R \tag{3.35}$$

corresponds via the adjunction $\coprod_\varepsilon \dashv \varepsilon^*$ to

$$\mathsf{Rel}(P)_{QX} \circ \mathsf{Rel}(Q)_X(R) \to \varepsilon^*(R) \tag{3.36}$$

This further corresponds via the definition of inverse image to the existence of a map $\varkappa$ in

$$\begin{tikzcd} \mathsf{Rel}(P)_{QX} \circ \mathsf{Rel}(Q)_X(R) \arrow[d, dashed] \arrow[dr, "\varkappa"] & \\ \varepsilon^*(R) \arrow[r, "\bar{\varepsilon}"'] & R \end{tikzcd} \tag{3.37}$$

This exists by the adjunction $\mathsf{Rel}(P) \dashv \mathsf{Rel}(Q)$ shown earlier. ■

The above lemma fulfils the first proof obligation stated in the beginning of Section 3.3.2, giving us the following picture:

$$\begin{tikzcd} \mathsf{Rel}_{QX} \arrow[loop left, "(\iota_X \circ Q\xi)^* \circ \mathsf{Rel}(B)_{QX}"] \arrow[r, bend left, "\coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX}"] \arrow[r, phantom, "\perp"] & \mathsf{Rel}_X \arrow[l, bend left, "\mathsf{Rel}(Q)_X"] \arrow[loop right, "\xi^* \circ \mathsf{Rel}(L)_X"] \end{tikzcd} \tag{3.38}$$

We now move on to the second proof obligation, i.e., that we have the equality/isomorphic step:

$$(\iota_X \circ Q\xi)^* \circ \mathsf{Rel}(B)_{QX} \circ \mathsf{Rel}(Q)_X = \mathsf{Rel}(Q)_X \circ \xi^* \circ \mathsf{Rel}(L)_X \tag{3.39}$$

To this end, we first show that $\mathsf{Rel}(Q)$ preserves inverse images, using the fact that we can obtain inverse images as pullbacks inside the category of relations. Since $\mathsf{Rel}(Q)$ is a right adjoint, it preserves these pullbacks.

Lemma 3.3.8: We obtain the inverse image as the pullback:

$$\begin{tikzcd} (f^*(S), X) \arrow[r, "\bar{f}(S)"] \arrow[d, tail] & (S, Y) \arrow[d, tail] \\ (\top, X) \arrow[r, "f \times f"] & (\top, Y) \end{tikzcd} \tag{3.40}$$

in the base category, where $\top$ is the largest relation, i.e., for an object X, the relation $X \times X$. ◂

Proof. This holds by definition of inverse image as a pullback in $\mathcal{C}$, as well as the definition of maps between relations. ■

Lemma 3.3.9: $\mathsf{Rel}(Q)$ preserves inverse images. ◂

Proof. As we have shown, $\mathsf{Rel}(Q)$ is a right adjoint, so it preserves limits. In particular, it preserves pullbacks, so that $\mathsf{Rel}(Q)$ applied to the pullback square obtained from Lemma 3.3.8 is again a pullback square

$$\begin{tikzcd} \mathsf{Rel}(Q)(f^*(S), X) \arrow[r, "\mathsf{Rel}(Q)(\bar{f}(S))"] \arrow[d, "\mathsf{Rel}(Q)(\mathrm{id}_X)"'] \arrow[dr, phantom, "\lrcorner", very near start] & \mathsf{Rel}(Q)(S, Y) \arrow[d, "\mathsf{Rel}(Q)(\mathrm{id}_Y)"] \\ \mathsf{Rel}(Q)(\top, X) \arrow[r, "\mathsf{Rel}(Q)(f)"] & \mathsf{Rel}(Q)(\top, Y) \end{tikzcd} \tag{3.41}$$

Now, as Q is a right adjoint, it preserves products so that $Q(\top_X) = \top_{QX}$. Thus, also $\mathsf{Rel}(Q)(\top, X) = (\top, QX)$. This means that the above diagram is of the same form as we have shown defines inverse images, meaning that $\mathsf{Rel}(Q)(f^*(S), X) = (Qf)^*(\mathsf{Rel}(Q)(S, Y))$. ■

We are now ready to show the existence of the required isomorphic step.

Theorem 3.3.10: If Q preserves abstract epis, then for any L-coalgebra (X, ξ):

$$(\iota_X \circ Q\xi)^* \circ \mathsf{Rel}(B)_{QX} \circ \mathsf{Rel}(Q)_X = \mathsf{Rel}(Q)_X \circ \xi^* \circ \mathsf{Rel}(L)_X \tag{3.42}$$ ◂

Proof. We have

$$\begin{aligned} & (\iota_X \circ Q\xi)^* \circ \mathsf{Rel}(B)_{QX} \circ \mathsf{Rel}(Q)_X \\ &= (Q\xi)^* \circ \iota_X^* \circ \mathsf{Rel}(B)_{QX} \circ \mathsf{Rel}(Q)_X && (3.43) \\ &= (Q\xi)^* \circ \mathsf{Rel}(Q)_{LX} \circ \mathsf{Rel}(L)_X && (3.44) \\ &= \mathsf{Rel}(Q)_X \circ \xi^* \circ \mathsf{Rel}(L)_X && (3.45) \end{aligned}$$

where Eq. (3.43) is an application of a basic fact on inverse images (technically, that the poset fibration of relations is split), Eq. (3.44) holds by Theorem 3.3.6, and Eq. (3.45) holds by Lemma 3.3.9. ■

We now reach our main result on preservation and reflection of bisimulations and bisimilarity by $\mathsf{Rel}(Q)_X$.

Theorem 3.3.11: Let (X, ξ) be an L-coalgebra. Suppose that Q preserves abstract epis. Then $\mathsf{Rel}(Q)_X$ maps bisimilarity on (X, ξ) to bisimilarity on $\overline{Q}(X, \xi)$. Further, $\mathsf{Rel}(Q)_X$ preserves bisimulations and, if it is order-reflecting, also reflects bisimulations. ◂

Proof. We have seen in Lemma 3.3.7 that $\mathsf{Rel}(Q)_X$ has a left adjoint, yielding the setting shown in Eq. (3.38). The equality in Theorem 3.3.10 then tells us that in this setting we have an isomorphic step. The result now follows from Lemma 3.3.2. ∎

While this result is formulated in terms of $\mathsf{Rel}(Q)_X$, we will also speak of $\overline{Q}$ (or even Q when the meaning is clear from context) preserving and reflecting both bisimulations and bisimilarity.

As a special case of Theorem 3.3.11, we recover (a version of) the following result found in [Rut00; BSV04; BSS17; BSS21].

Lemma 3.3.12: Assume functors $B, L \colon \mathcal{C} \to \mathcal{C}$, and a natural transformation $\iota \colon L \to B$. Then the functor $\overline{\mathsf{Id}} \colon \mathsf{Coalg}(L) \to \mathsf{Coalg}(B)$ defined by $(X, \xi) \mapsto (X, \iota_X \circ \xi)$ on objects and identity on morphisms, preserves bisimulations. If additionally ι has a left inverse, $\overline{\mathsf{Id}}$ reflects bisimulations. ◂

Proof. Viewing the transformation ι as a backward step for the adjunction $\mathsf{Id} \dashv \mathsf{Id} \colon \mathcal{C} \to \mathcal{C}$ gives us preservation by $\overline{\mathsf{Id}}$. If ι has a left inverse, we have an invertible step. In this case the right adjoint trivially preserves abstract epis. Also, $\mathsf{Rel}(\mathsf{Id})_X = \mathsf{Id}_{\mathsf{Rel}_X}$ is clearly order-reflecting, so we have preservation and reflection of bisimulations by $\mathsf{Rel}(\mathsf{Id})_X$ and the lifting $\overline{\mathsf{Rel}(\mathsf{Id})_X}$ preserves and reflects bisimilarity. ∎

A mate on the fibres

Our results so far show when the right adjoint in the setting of step-induced liftings preserves and reflects bisimulations and bisimilarity. In the process, we have shown that the relation lifting $\mathsf{Rel}(Q)$ restricted to fibres has a left adjoint. It turns out that (one half of) the equality of Theorem 3.3.10 also entails preservation of bisimulations by the left adjoint $\coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX}$ of Lemma 3.3.7. The main idea is that the forward step part of the equality has a mate, which allows us to lift the left adjoint to post-fixed points. Note that this does not require the property that Q preserves abstract epis used for the proof of the full equality, as we show now.

Proposition 3.3.13: Let (X, ξ) be an L-coalgebra. The map $\coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX}$ preserves bisimulations, i.e., it sends a bisimulation on $\overline{Q}(X, \xi)$ to a bisimulation on (X, ξ). ◂

Proof. By the first part of Proposition 3.3.5, δ lifts to relations. As shown in Theorem 3.3.6, this means that we have $\iota_X^* \circ \mathsf{Rel}(B)_{QX} \circ \mathsf{Rel}(Q)_X \leq \mathsf{Rel}(Q)_{LX} \circ \mathsf{Rel}(L)_X$. By the same reasoning as in Theorem 3.3.10, we then have $(\iota_X \circ Q\xi)^* \circ \mathsf{Rel}(B)_{QX} \circ \mathsf{Rel}(Q)_X \leq \mathsf{Rel}(Q)_X \circ \xi^* \circ \mathsf{Rel}(L)_X$.

The mate of this inequality is then $\coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX} \circ (\iota_X \circ Q\xi)^* \circ \mathsf{Rel}(B)_{QX} \leq \xi^* \circ \mathsf{Rel}(L)_X \circ \coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX}$, so that if $R \leq (\iota_X \circ Q\xi)^* \circ \mathsf{Rel}(B)_{QX}(R)$ (i.e., R is a bisimulation on $\overline{Q}(X, \xi)$), then

$$\coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX}(R) \leq \coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX} \circ (\iota_X \circ Q\xi)^* \circ \mathsf{Rel}(B)_{QX}(R) \tag{3.46}$$

$$\leq \xi^* \circ \mathsf{Rel}(L)_X \circ \coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX}(R) \tag{3.47}$$

meaning $\coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX}(R)$ is a bisimulation on (X, ξ). ∎

Order-reflection

We briefly turn to the condition of order-reflection. As we are often interested in cases where the right adjoint is a forgetful functor in the context of an Eilenberg-Moore adjunction, we will show the condition holds for that case. For this, we will require the following intermediate result to later conclude that maps between relations on algebras above the identity are themselves algebra morphisms.

Lemma 3.3.14: For a monad $\mathcal{T}$ on $\mathcal{C}$, let $(X, x), (Y, y), (Z, z)$ be objects in $\mathcal{EM}(\mathcal{T})$, and $f : (X, x) \to (Z, z)$, $h : (Y, y) \rightarrowtail (Z, z)$ algebra morphisms with h mono. Further, suppose we have $g : X \to Y$ a morphism in $\mathcal{C}$, such that $h \circ g = f$. Then g is in fact a morphism in $\mathcal{EM}(\mathcal{T})$. ◂

Proof. Consider

$$\begin{tikzcd} TX \arrow[r, "Tg"] \arrow[d, "x"] \arrow[rr, bend left, "Tf"] & TY \arrow[r, "Th"] \arrow[d, "y"] & TZ \arrow[d, "z"] \\ X \arrow[r, "g"] \arrow[rr, bend right, "f"'] & Y \arrow[r, tail, "h"] & Z \end{tikzcd} \tag{3.48}$$

We need to show that $g \circ x = y \circ Tg$. As h is mono, it suffices to show that $h \circ g \circ x = h \circ y \circ Tg$. Indeed, we have

$$h \circ g \circ x = f \circ x \tag{3.49}$$
$$= z \circ Tf \tag{3.50}$$
$$= z \circ Th \circ Tg \tag{3.51}$$
$$= h \circ y \circ Tg \tag{3.52}$$

completing the proof. ■

Lemma 3.3.15: For a monad $\mathcal{T}$ with forgetful functor $U\colon\ \mathcal{EM}(\mathcal{T}) \to \mathcal{C}$, the (restricted) lifting $\mathsf{Rel}(U)_X$ is an order-reflecting map. ◂

Proof. Suppose we have two relations $R, S \in \mathsf{Rel}_X$, with $\mathsf{Rel}(U)(R) \leq \mathsf{Rel}(U)(S)$. We have the following commutative diagram

$$\begin{tikzcd} \mathsf{Rel}(U)(R) \arrow[rr, "f"] \arrow[dr, tail, "r"'] & & \mathsf{Rel}(U)(S) \arrow[dl, tail, "s"] \\ & UX \times UX & \end{tikzcd} \tag{3.53}$$

Then by Lemma 3.3.14, f is also an algebra morphism so that $R \leq S$ in Rel_X. ■

If $\mathcal{C} = \mathsf{Set}$ in the above lemma, then $\mathsf{Rel}(U)_X$ is just the inclusion of the poset of congruences $\mathsf{Rel}_{(X,\alpha)}$ on an algebra (X, α) into the poset of all relations on its carrier.

In that case, we can also use the above to show preservation and reflection of *behavioural equivalence*. Two states of a coalgebra (in Set) are behaviourally equivalent if they can be identified by some coalgebra homomorphism. This can be captured more abstractly using kernel bisimulations (see, e.g., [Sta11]). Since U is assumed to be a forgetful functor to Set, we simply define preservation and reflection of behavioural equivalence by $\overline{U}$ to mean that for any two states x, y of an L-coalgebra (X, ξ), x and y are behaviourally equivalent (for (X, ξ)) if and only if they are behaviourally equivalent for $\overline{U}(X, \xi)$.

It turns out that, in our setting, coincidence of bisimilarity and behavioural equivalence for L-coalgebras reduces to coincidence for B-coalgebras. This is stated in the following lemma; the essence is that $\overline{U}$ is easily shown to preserve behavioural equivalence.

Lemma 3.3.16: For a monad $\mathcal{T}$, consider the Eilenberg-Moore adjunction $\mathcal{F} \dashv U\colon\ \mathcal{EM}(\mathcal{T}) \to \mathsf{Set}$. Let $L\colon\ \mathcal{EM}(\mathcal{T}) \to \mathcal{EM}(\mathcal{T})$ and $B\colon\ \mathsf{Set} \to \mathsf{Set}$ be functors, and $\delta\colon\ BU \to UL$ an invertible step. Further suppose that $\overline{U}$ reflects bisimilarity, and that B preserves weak pullbacks. Then bisimilarity and behavioural equivalence for L-coalgebras coincide (and hence, $\overline{U}$ preserves and reflects behavioural equivalence). ◂

Proof. Bisimilarity always implies behavioural equivalence [Sta11]. Conversely, suppose that we have two states $x, y \in X$ of a coalgebra $\xi\colon X \to LX$ in $\mathcal{EM}(\mathcal{T})$, which are behaviourally equivalent. We have some coalgebra $(Z, \zeta\colon Z \to LZ)$ and coalgebra homomorphisms $g, h\colon (X, \xi) \to (Z, \zeta)$, such that $g(x) = h(y)$. Then in $\overline{U}(X, \xi)$, the states x, y will be identified in $\overline{U}(Z, \zeta)$ by $\overline{U}(g)$ and $\overline{U}(h)$ as $\overline{U}$ acts as the identity on morphisms. Therefore, x, y are behaviourally equivalent in $\overline{U}(X, \xi)$. Since B preserves weak pullbacks, this means that x, y are bisimilar [Sta11]. Finally, as we have assumed that bisimilarity is reflected by $\overline{U}$, the states x, y are bisimilar in the original coalgebra ξ. ■

3.4 Applications

Now that we have obtained conditions for the preservation and reflection of bisimilarity, we return to the examples presented in Section 3.2. We will show how a number of existing non-trivial results can be recovered using our conditions. Further, the Set-Stone adjunction used in the first example is known to not be monadic[2], and so outside the scope of weak liftings, which indicates the generality of our results.

3.4.1 Ultrafilter Extensions and Vietoris bisimulations

In Section 3.2.2, we have seen how the construction of ultrafilter extensions (of Kripke frames) can be obtained from an invertible step, which arises from a weak lifting described by Garner. In the current treatment of reflection and preservation of bisimilarity, we focus on the restriction of this invertible step to the category Stone:

$$\begin{tikzcd} \mathsf{Set} \arrow[loop left, "\mathcal{P}"] \arrow[r, bend left, "\mathcal{F}"] \arrow[r, phantom, "\perp"] & \mathsf{Stone} \arrow[l, bend left, "U"] \arrow[loop right, "\mathcal{V}"] \end{tikzcd} \qquad \mathcal{P}U \xrightarrow{\delta} U\mathcal{V} \xrightarrow{\iota} \mathcal{P}U \tag{3.54}$$

This brings us in line with [BFV10], where a comparison is made between bisimilarity for the Vietoris functor $\mathcal{V}\colon$ Stone $\to$ Stone and bisimilarity for

[2]This seems to be folklore. It follows, for example, from the fact that Stone is not exact [Bar71], and the result (see, e.g., [Bor94b, Thm. 4.4.5]) that categories monadic over Set are exact. Alternatively, it holds because Stone is a proper subcategory of CHaus and the monads induced by the adjunctions with Set are the same.

the powerset functor $\mathcal{P}\colon \mathsf{Set} \to \mathsf{Set}$, called Vietoris-bisimilarity and Kripke-bisimilarity respectively in *op. cit.* More precisely, for a $\mathcal{V}$-coalgebra $(\boldsymbol{X}, \gamma)$, Kripke bisimilarity is bisimilarity on $\overline{U}(\boldsymbol{X}, \gamma)$, where $\overline{U}$ is the step-induced lifting of the forgetful functor $U\colon \mathsf{Stone} \to \mathsf{Set}$. Vietoris bisimilarity is simply bisimilarity on the coalgebra $(\boldsymbol{X}, \gamma)$ itself.

Applying the framework of Section 3.3 to the above setting, we recover a number of results from [BFV10].

Proposition 3.4.1 ([BFV10, Prop 3.4]): The relation liftings of $\mathcal{P}\colon \mathsf{Set} \to \mathsf{Set}$ and $\mathcal{V}\colon \mathsf{Stone} \to \mathsf{Stone}$ coincide for closed relations. Precisely, for any $R \subseteq \boldsymbol{X} \times \boldsymbol{X}$ closed in the product topology, we have

$$\mathsf{Rel}(U) \circ \mathsf{Rel}(\mathcal{V})(R) = (\mathsf{Rel}(\mathcal{P}) \circ \mathsf{Rel}(U)(R)) \cap (U\mathcal{V}\boldsymbol{X} \times U\mathcal{V}\boldsymbol{X}) \qquad ◀$$

Proof. We can apply Theorem 3.3.6, as U preserves abstract epis by the same argument as for adjunctions monadic over Set (see the discussion after Proposition 3.3.5), as Stone is also a regular category. This gives the equality $\mathsf{Rel}(U)_{\mathcal{V}\boldsymbol{X}} \circ \mathsf{Rel}(\mathcal{V})_{\boldsymbol{X}} = \iota^*_{\boldsymbol{X}} \circ \mathsf{Rel}(\mathcal{P})_{U\boldsymbol{X}} \circ \mathsf{Rel}(U)_{\boldsymbol{X}}$, where the action of $\iota^*_{\boldsymbol{X}}\colon \mathsf{Rel}_{\mathcal{P}U\boldsymbol{X}} \to \mathsf{Rel}_{U\mathcal{V}\boldsymbol{X}}$ is precisely the restriction to closed subsets. ■

Theorem 3.4.2 ([BFV10, Thm 3.6]): Vietoris bisimulations are equivalently closed Kripke bisimulations. ◀

Proof. For this, we use preservation and reflection of bisimulations by the (restricted) relation lifting $\mathsf{Rel}(U)_{\boldsymbol{X}}\colon \mathsf{Rel}_{\boldsymbol{X}} \to \mathsf{Rel}_{U\boldsymbol{X}}$ of the forgetful functor, which is simply the inclusion of the poset of closed relations on a Stone space $\boldsymbol{X}$ to that of all relations on the underlying set.

Indeed, preservation and reflection by $\mathsf{Rel}(U)_{\boldsymbol{X}}$ follows from Theorem 3.3.11. We have seen that $U\colon \mathsf{Stone} \to \mathsf{Set}$ preserves abstract epis, so it only remains to check that $\mathsf{Rel}(U)_{\boldsymbol{X}}$ is order-reflecting. This holds because Stone is a full subcategory of CHaus, which is monadic over Set, so that the discussion of the previous section applies. ■

Theorem 3.4.3 ([BFV10, Thm 5.2]): The closure of a Kripke bisimulation is a Vietoris bisimulation. ◀

Proof. The left adjoint $\coprod_\varepsilon \circ \mathsf{Rel}(\mathcal{F})_{U\boldsymbol{X}}$ of Lemma 3.3.7 gives the closure of a relation. The inclusion $\coprod_\varepsilon \circ \mathsf{Rel}(\mathcal{F})_{U\boldsymbol{X}}(R) \subseteq \overline{R}$ can be seen to hold by the equivalence with the inclusion $R \subseteq \mathsf{Rel}(U)_{\boldsymbol{X}}(\overline{R})$ via the adjunction. For the other direction, we know that $\coprod_\varepsilon \circ \mathsf{Rel}(\mathcal{F})_{U\boldsymbol{X}}(R)$ is closed as it is a relation on a Stone space. It further contains R as $\mathsf{Rel}(\mathcal{F})_{U\boldsymbol{X}}(R)$ contains all pairs of principal ultrafilters generated by x, y for $(x, y) \in R$. The counit acts by ultrafilter convergence, so that the direct image recovers elements of R. By definition of the closure $\overline{R}$ as the smallest closed set containing R we then have $\overline{R} \subseteq \coprod_{\varepsilon} \circ \mathsf{Rel}(\mathcal{F})_{UX}(R)$ The result thus follows from Proposition 3.3.13. ■

Theorem 3.4.4 ([BFV10, Cor 3.10]): Let $\gamma : \boldsymbol{X} \to \mathcal{V}\boldsymbol{X}$ be a coalgebra and $x, y \in X$ states. Then x and y are bisimilar in $(\boldsymbol{X}, \gamma)$ iff they are bisimilar in $\overline{U}(\boldsymbol{X}, \gamma)$. I.e., Vietoris- and Kripke-bisimilarity are equivalent. ◂

Proof. It follows from Theorem 3.3.11 that $\mathsf{Rel}(U)_{\boldsymbol{X}}$ maps the greatest fixed point, say $\sim_{\mathcal{V}}$, of $\gamma^* \circ \mathsf{Rel}(\mathcal{V})_{\boldsymbol{X}}$ to that of $(\iota_{\boldsymbol{X}} \circ U\gamma)^* \circ \mathsf{Rel}(\mathcal{P})_{U\boldsymbol{X}}$ (call this $\sim_{\mathcal{P}}$). The action of $\mathsf{Rel}(U)_{\boldsymbol{X}}$ is then such that pairs are preserved so that $x \sim_{\mathcal{V}} y$ iff $x \sim_{\mathcal{P}} y$. ■

3.4.2 PAs and Belief State Transformers

As discussed in Section 3.2.2, we can determinise a PA to a coalgebra for the convex powerset functor $\mathcal{P}_c : \mathcal{EM}(\mathcal{D}) \to \mathcal{EM}(\mathcal{D})$ using a lifting of $\mathcal{F} : \mathsf{Set} \to \mathcal{EM}(\mathcal{D})$. The step-induced lifting of the corresponding forgetful functor $U : \mathcal{EM}(\mathcal{D}) \to \mathsf{Set}$ maps the $\mathcal{P}_c$-coalgebra back into Set, but we must take care that this does not change its behaviour.

What we will do now is use the results of Section 3.3 to show that bisimilarity is preserved and reflected. Once we know this, we further apply Lemma 3.3.16 to show the coincidence of bisimilarity and behavioural equivalence in the case of the convex powerset functor on $\mathcal{EM}(\mathcal{D})$ and the powerset functor on Set. This coincidence is relevant for the generalisation of the corresponding results of [BSS21] (restricted to the convex powerset functor), which are formulated in terms of behavioural equivalence. We thus recover the following result:

Proposition 3.4.5 ([BSS21, Proposition 6.6]): The step-induced lifting of the forgetful functor $U : \mathcal{EM}(\mathcal{D}) \to \mathsf{Set}$ preserves and reflects behavioural equivalence on $\mathcal{P}_c^{\Sigma}$-coalgebras. ◂

Proof. We have seen that U preserves abstract epis as the adjunction in question is monadic over Set. This allows us to apply Theorem 3.3.11 so that $\overline{U}$ indeed preserves and reflects bisimulations, and the relevant lifting preserves and reflects bisimilarity. From Lemma 3.3.16, it follows that $\overline{U}$ also preserves and reflects behavioural equivalence. ■

We also recover [BSS21, Proposition 6.5], which relates kernel bisimulations for $\mathcal{P}_c^{\Sigma}$-coalgebras and their image under the lifted forgetful functor.

Proposition 3.4.6: A relation R is a kernel bisimulation for a $\mathcal{P}_c^{\Sigma}$-coalgebra $(\mathbb{S}, \gamma)$ in $\mathcal{EM}(\mathcal{D})$ iff it is a kernel bisimulation for $\overline{U}(\mathbb{S}, \gamma)$ and also a congruence. ◂

Proof. Assuming $(\mathbb{S}, \gamma)$ is a $\mathcal{P}_c^\Sigma$-coalgebra, this follows from Lemma 3.3.16 together with the previous item, and the fact that bisimulations in Eilenberg-Moore categories are congruences. ∎

In [BSS21], the correspondence is proved in the setting of a *lax lifting*, rather than the weak lifting we use. In our notation, L is a lax lifting of B if there is an *injective* natural transformation $\iota \colon QL \to BQ$. This is close to a weak lifting: components of such a ι for the above functors are injective functions and each domain QLX will be non-empty as $L = \mathcal{P}_c^\Sigma$. The components thus have left inverses. However, these inverses may not form a *natural* transformation. See Chapter 4 for a further discussion of naturality of weak liftings.

3.5 Conclusion

We studied the notion of an *invertible step*, which provides several constructions on coalgebras via functor liftings. We showed that the lifting of the right adjoint, induced by such an invertible step, preserves and reflects bisimilarity. This abstract result instantiates to several concrete results from the literature, in examples related to ultrafilter extensions and weak distributive laws.

We have focused on preservation and reflection of bisimilarity, defined in terms of relation lifting. There are several other coalgebraic notions of behavioural equivalence and bisimilarity [Sta11]—we discuss these in the next subsection. Finally, in Section 3.5.2 we list directions for future work.

3.5.1 Remarks on other notions of bisimulation

Aczel-Mendler bisimulations

An alternative notion of bisimulation for a coalgebra $\gamma \colon X \to LX$, are Aczel-Mendler bisimulations. These are defined as relations $R \rightarrowtail X \times X$ with an L-coalgebra structure $\gamma_R \colon R \to LR$ on R such that the projection maps into X are coalgebra homomorphisms from (R, γ_R) to (X, γ) [AM89].

One may try to apply a similar approach as in the preceding chapters to show preservation and reflection of such bisimulations, however this encounters a number of problems, which we sketch here. The setting of ultrafilter extensions also gives an example where Aczel-Mendler bisimulations do not yield the desired notion of equivalence, for which the classical "bisimilarity-somewhere-else" result holds.

To use our earlier techniques to speak of preservation of Aczel-Mendler bisimulations, we would apply a lifting $\overline{Q}$ to such a bisimulation R. This yields a structure $QR \to BQR$, however, this is not immediately a bisimulation, as QR may not be a relation. We can obtain a relation by taking the image of $\langle Q\pi_1, Q\pi_2\rangle$ as we do to define relation lifting, but in general this is a coalgebraic (Hermida-Jacobs) bisimulation [Jac16, Exercise 4.5.2], rather than an Aczel-Mendler one.

To instead speak of reflection of Aczel-Mendler bisimulations, we start with a span $QX \leftarrow R \rightarrow QX$ and try to construct a relation on X. Using the adjunction of the step setting, we can transpose the projections to obtain a span $X \leftarrow PR \rightarrow X$. Again PR is not immediately a relation in general, and taking the image yields a $\mathsf{Rel}(L)$-coalgebra (not an L-coalgebra) as the projections and the counit ε are coalgebra homomorphisms (see also [Jac16, Exercise 4.5.4]). This in fact comes down to the same as the left adjoint $\coprod_\varepsilon \circ \mathsf{Rel}(P)_{QX}$ constructed earlier. There we factorise to obtain the relation lifting and factorise again for the direct image of ε, instead of factorising the paired transposes defined using ε. We also do not explicitly use that ε is a coalgebra homomorphism (although this follows from the step with right inverse and Lemma 3.2.10); instead we lift the adjunction at the level of relations to give a map between bisimulations. This is part of the motivation for the use of relation liftings and the corresponding notion of bisimulations.

Going further, it is shown in [BFV10] that there exists a Vietoris bisimulation which is not an Aczel-Mendler bisimulation and, stronger, that there exist Vietoris coalgebras with states which can be related by a Vietoris but not an Aczel-Mendler bisimulation. Thus, the correspondences between bisimulations on Set and Stone we have discussed in the previous sections are not obtainable when we consider Aczel-Mendler bisimulations.

Kernel bisimulations/behavioural equivalence

In applying our results to the preservation and reflection of behavioural equivalence, we currently work concretely; considering sets of states and identification of elements.

A more abstract approach may be to consider kernel bisimulations. A relation $R \rightarrowtail X \times X$ is a kernel bisimulation on a coalgebra $(X, \gamma\colon X \to LX)$ in a category $\mathcal{D}$, if it is the pullback of morphisms $X \to Z \leftarrow X$ in $\mathcal{D}$ forming a cospan of coalgebra homomorphisms $(X,\gamma) \to (Z,\zeta) \leftarrow (X,\gamma)$ in $\mathrm{Coalg}_{\mathcal{D}}(L)$. In a concrete setting this coincides with behavioural equivalence, as such a pullback contains exactly the pairs of elements of X which are identified in Z by the morphisms forming the cospan. We can thus view this as a generalisation of behavioural equivalence as defined earlier.

Assuming an invertible step $\delta\colon BQ \to QL$, we would like to relate R to a kernel bisimulation on the coalgebra $\overline{Q}(X,\gamma)$ obtained by applying the step-induced lifting of Q. Applying Q to the pullback square for R yields a pullback square as Q is a right adjoint. However, as in our discussion of Aczel-Mendler bisimulations, this may not be a relation. We may try to also use relation liftings here, and take $\mathsf{Rel}(Q)(R)$ instead of $Q(R)$, however this may no longer be a pullback. It is not currently clear to us how to resolve these problems in general.

3.5.2 Looking forward

As suggested in the paper on which this chapter is based, in the coming chapter we go beyond preservation and reflection of bisimilarity, and consider more general *(co)inductive predicates* in fibrations (see, e.g., [HKC18]). We also move from invertible steps to studying *forward* and *backward* steps separately, i.e., we do not require them to be each other's one-sided inverses. This allows the development of separate conditions for preservation and reflection by left and right adjoints.

This separation will further allow application to examples beyond the Eilenberg-Moore (and similar) adjunctions we have considered so far, covering among others conditions for expressivity of coalgebraic modal logics.

Some remaining open questions lie in Section 3.4 and the application of our results to the work of [BFV10]. This has been generalised in two directions: the recent [GT22] considers bisimulations for Vietoris coalgebras on the category of *arbitrary* topological spaces, while [ES18] develops a notion of neighbourhood bisimulation for coalgebras that allows to generalise the results from [BFV10] to a large variety of functors on the category of Stone spaces and their corresponding functors on Set. We would like to understand whether our framework is able to recover these generalisations.

CHAPTER 4

Forward and Backward Steps in a Fibration

$B^{\Delta 7}$ D^7 $G^{\Delta 7}$ $B\flat^7$ $E\flat^{\Delta 7}$ Am^7 D^7
$G^{\Delta 7}$ $B\flat^7$ $E\flat^{\Delta 7}$ $F\sharp^7$ $B^{\Delta 7}$ Fm^7 $B\flat^7$
$E\flat^{\Delta 7}$ Am^7 D^7 $G^{\Delta 7}$ $C\sharp m^7$ $F\sharp^7$
$B^{\Delta 7}$ Fm^7 $B\flat^7$ $E\flat^{\Delta 7}$ $C\sharp m^7$ $F\sharp^7$

John Coltrane

4.1 Introduction

We continue the discussion of the topics of the previous chapter with a more general look at transformations between systems modelled as coalgebras for various functors. More specifically, transformations which can be defined as liftings of adjoint functors.

Such liftings are central to, for instance, coalgebraic approaches to trace semantics and determinisation [HJS07; JSS15; Sil+13; BMS13; KKW14; RJL21] as well as testing semantics and algebraic semantics of modal logics [KKP04; PMW06; BK05; Kli07; CJ14; Kli05].

Our central question remains how these constructions on coalgebras affect notions of equivalence. For instance, determinisation of automata turns a coalgebra on, e.g., the category Set, into a coalgebra on the category of Eilenberg-Moore algebras for a monad, so that the canonical notion of behavioural equivalence changes from bisimilarity to language semantics. Subsequently, the algebraic structure may be forgotten again, turning the determinised coalgebra back into a Set coalgebra, and this simple operation does not affect behavioural equivalence. But when determinisation is not given by a distributive law, such as in the construction of belief-state transformers in [BSS21], proving that this “forgetting” preserves and reflects behavioural equivalence can be non-trivial (see *op. cit.*, and the previous chapter).

A different type of example of a construction given by a lifting of an adjoint, which is not covered by the results of the previous chapter, is the algebraic semantics of modal logic. A standard approach is based on a dual adjunction [KKP04; Gro22], for which a lifting of the left adjoint transforms coalgebras for a behaviour functor of interest (e.g., the powerset functor modelling the transitions of a Kripke model) into algebras for the functor whose initial algebra consists of the formulas of the modal logic. As we show later, one form of preservation of behavioural equivalence amounts to adequacy and expressivity of the logic. In fact, we obtain conditions for adequacy and expressivity without requiring an initial algebra for the modal logic functor, by instead relating definability of predicates to bisimilarity.

We propose an abstract framework to analyse whether a coalgebra lifting of an adjoint preserves various properties of coalgebras. The properties we consider in this chapter, are those definable as (co)inductive invariants and predicates [HKC18]. These are defined using functor liftings in fibrations, first used to define a coalgebraic version of bisimilarity [HJ98]. It has more recently been shown that coinductive predicates can be used to give general proofs of expressivity of modal logics [KR21; Kom+21].

The coalgebra liftings of adjoints will be those arising from natural transformations recently termed *steps* [RJL21]. Also sometimes called *morphisms of endofunctors*, steps are a variant of distributive laws involving a left or right adjoint. It has been shown that they cover the above-mentioned examples of language semantics, determinisations, and modal logic.

The goal of this chapter will thus be to connect steps and fibrations; to speak generally about preservation of coinductive (and inductive) predicates by transformations of coalgebras. The key technical idea is to use a variant of fibred adjunctions [Jac01]. We start with an adjunction and a step, and assume a fibration and functor lifting on both sides of the adjunction to formulate the coinductive predicates that we wish to relate. We then lift the underlying adjunction to the total categories of the fibrations involved [Her93]. These liftings then restrict to the fibres in which we can formulate sufficient conditions for preservation of coinductive predicates by coalgebra constructions induced by steps.

There are two main variants of this abstract story: one that starts from a step that lifts the *left* adjoint to coalgebras, and one for lifting the *right* adjoint. The first allows us, for instance, to speak about adequacy and expressivity of modal logics, without referring to initial algebras. This connects to recent work that uses Galois connections [Beo+23], and in fact we recover those Galois connections from our adjunctions between fibrations. We also study an example that has not occurred in previous abstract frameworks for expressivity: proving expressivity of a logic by relating it to *apartness* [GJ21; Geu22] instead of bisimilarity.

The latter constructions, arising as liftings of right adjoints, include part of the construction of belief-state transformers from probabilistic automata, as studied in [BSS21]. More precisely, it follows from our results that any right adjoint in a *lax lifting* situation preserves and reflects bisimilarity (with the usual assumption of injectivity replaced by assuming split monomorphisms), generalising the result for belief-state transformers.

While our focus is on (co)inductive *predicates* and the instances of bisimilarity and modal definability/logical equivalence, we will also briefly discuss how the conditions which can be formulated in the fibrational setting relate to preservation and reflection of (co)inductive *invariants*. This relates to the conditions for preservation and reflection of bisimulations of the previous chapter. For instance, we are able to show that the convex closure of (a certain class of) bisimulations are again bisimulations. A similar result was shown earlier for topological closure.

Contributions Our main contributions in this chapter are: first, conditions relating steps and functor liftings ensuring preservation of coinductive predicates (Corollaries 4.4.5 and 4.4.9); and second similar conditions implying preservation of coinductive invariants (Lemmata 4.5.1 and 4.5.2). Similarly to the last chapter, these results show certain forms of preservation and reflection of fixpoints by liftings of adjoint functors, but go beyond the earlier results thanks to the more general setting of $\mathsf{CLat}_\sqcap$ fibrations.

For the application of these results we provide constructions of lifted adjunctions to the fibrations of (equivalence) relations, predicates, and quotients. This allows us to recover results in the settings of: lax liftings where notions of equivalence can be compared in fibrations of relations; and expressivity of coalgebraic logic where the logical semantics and bisimilarity can be defined and related as predicates and relations.

Finally, we note that the conditions for expressivity of a coalgebraic modal logic are obtained both by comparing logical equivalence and bisimilarity as usual, but also by comparing logical equivalence with apartness, the dual of bisimilarity. This dualisation of notions of equivalence will also be the focus of the second half of this thesis.

4.2 Preservation of coinductive predicates in lattices

Before moving to the general theory of fibrations and steps, we consider some examples of the special case of steps between lattices assumed in Lemma 3.3.2. In that order-theoretic setting, steps are simply inequalities between compositions of monotone maps (an instance of natural transformations between composed functors), and their existence implies preservation and reflection of post-fixed points and greatest fixed points in the sense of Definition 3.3.1.

At the end of the chapter, we will return to the examples presented here to show how they arise from liftings of adjunctions and what preservation and reflection of various notions of equivalence means concretely in these contexts.

4.2.1 Example: Closed and Convex Relations

We will first consider two instances where the lattices consist of relations on sets on one side, and relations with either topological or convex structure on the other, i.e.:

$$\begin{tikzcd} \mathsf{Rel}_X \arrow[loop left, "b"] \arrow[r, bend left, "c"] \arrow[r, phantom, "\perp"] & \mathsf{CRel}_{\boldsymbol{X}} \arrow[l, bend left, "u"] \arrow[loop right, "v"] \\ \mathsf{Rel}_A \arrow[loop left, "b"] \arrow[r, bend left, "h"] \arrow[r, phantom, "\perp"] & \mathsf{ConRel}_{\boldsymbol{A}} \arrow[l, bend left, "u"] \arrow[loop right, "d"] \end{tikzcd} \tag{4.1}$$

where $\mathsf{CRel}_{\boldsymbol{X}}$ consists of closed relations on a compact Hausdorff space $\boldsymbol{X}$ and $\mathsf{ConRel}_{\boldsymbol{A}}$ consists of convex relations on a convex algebra $\boldsymbol{A}$. The monotone maps v, d will be such that the post-fixed points are bisimulations and the greatest fixed points are bisimilarity on systems with topological ($\boldsymbol{X}$) or convex ($\boldsymbol{A}$) structure, and the maps b on both sides are such that we obtain bisimulations/bisimilarity on systems where this structure has been forgotten (X and A respectively). Lemma 3.3.2 thus tells us how bisimilarity on each side can be related via the right adjoint.

These settings arise in examples of ultrafilter extensions for coalgebras and the transformation of probabilistic automata into belief-state transformers. In the first instance, the closed relations can in fact be restricted to Stone topological spaces (those compact Hausdorff spaces which are zero-dimensional), where we consider coalgebras for the Vietoris functor on Stone. It is shown in [BFV10] that these coalgebras correspond to descriptive frames, which arise in the first stage of the construction of the ultrafilter extension of a

Kripke frame given in [KKP05]. The second stage given there is to transport these back to a coalgebra in Set.

The construction of a belief-state transformer from a probabilistic automaton has a similar structure, where the second stage is to transport a system with extra algebraic structure back to a Set coalgebra. In each case, it is important that behavioural equivalence is preserved and reflected in the second stage, shown in [BFV10; BSS21] respectively for the above examples.

The results of [BFV10] were recovered already in the previous chapter. Here, we generalise our earlier approach so as to be directly applicable to the *lax* lifting setting of [BSS21], and to adequacy and expressivity of modal logics, which we show by relating logical equivalence to bisimilarity (as in the coming example) and apartness (Section 4.6.3).

4.2.2 Example: Expressivity

Another example relates to work on expressivity of coalgebraic logics [KR21; Beo+23; JS10; Kli07; Kli05], where we wish to relate bisimilarity and logical equivalence (or indistinguishability). The lattices we use here are those of equivalence relations on the carrier X of a coalgebra and predicates on 2^X. This gives us the following setting:

$$\begin{tikzcd} \mathsf{ERel}_X \arrow[loop left, "b"] \arrow[r, bend left, "f"] \arrow[r, phantom, "\perp"] & \mathsf{Pred}_{2^X}^{\mathrm{op}} \arrow[l, bend left, "g"] \arrow[loop right, "l"] \end{tikzcd} \tag{4.2}$$

where we define

$$g(\Gamma \subseteq 2^X) = \{(x, y) \mid \forall S \in \Gamma.\, x \in S \iff y \in S\} \tag{4.3}$$

In words, we relate those elements which are in exactly the same sets of Γ. Next, the action of f on an equivalence relation R is to give those subsets which are closed under R, i.e., they are a union of equivalence classes of R. These may also be called $\mathcal{R}$-coherent sets. More formally, we have

$$f(R) = \{\Gamma \subseteq X \mid \forall (x, y) \in R.\, x \in \Gamma \iff y \in \Gamma\} \tag{4.4}$$

The monotone map b is taken to be such that the greatest fixed point is bisimilarity. The map l is dual to the map whose *least* fixed point we can think of as those predicates obtainable as the interpretation of a formula of a modal logic, also called the *definable subsets*. In essence, these *are* the formulas which we generate from some propositional constants and applications of the operators of our logic.

Applying g to these "reachable" predicates gives an equivalence relating states which satisfy exactly the same formulas. This is exactly logical equivalence, and the above picture then allows us to relate this to bisimilarity. Namely, if $g(\nu l) \geq \nu b$, then bisimilarity implies logical equivalence, which is precisely adequacy of the logic. If, conversely, $g(\nu l) \leq \nu b$, then logical equivalence implies bisimilarity, called expressivity of the logic.

Now, Lemma 3.3.2 gives us a way to obtain these inclusions via inequalities capturing the interaction between the logic and behaviour in a rather general way. Later, we will show in more detail how these conditions relate to existing approaches to the semantics of coalgebraic modal logic and the properties of adequacy and expressivity.

4.3 Lifting adjunctions in a fibration

The goal of this section is to provide a general construction of the adjunctions required to compare (co)inductive predicates. To this end, we first show how to lift general adjunctions to the total categories of some classes of fibrations, in which many coinductive predicates of interest exist. After this, we show how such lifted fibrations can (in general) be restricted to the fibres in which we define those predicates. The approach is inspired by the results regarding lifted/fibred adjunctions of Hermida and Jacobs in [Her92; Her93; Jac01].

We start by defining what it means for an adjunction to lift to the total categories of fibrations.

Definition 4.3.1: Let $F \dashv G \colon \mathcal{D} \to \mathcal{C}$ be an adjunction, and assume fibrations $p \colon \mathcal{E} \to \mathcal{C}$ and $q \colon \mathcal{F} \to \mathcal{D}$. Further, assume an adjunction $\overline{F} \dashv \overline{G} \colon \mathcal{F} \to \mathcal{E}$. If we have the following commutative diagram

$$\begin{tikzcd} \mathcal{E} \arrow[r, bend left, "\overline{F}"] \arrow[r, phantom, "\perp"] \arrow[d, "p"'] & \mathcal{F} \arrow[l, bend left, "\overline{G}"] \arrow[d, "q"] \\ \mathcal{C} \arrow[r, bend left, "F"] \arrow[r, phantom, "\perp"] & \mathcal{D} \arrow[l, bend left, "G"] \end{tikzcd} \tag{4.5}$$

and the unit and counit of the adjunction $\overline{F} \dashv \overline{G}$ are above the unit and counit of the adjunction $F \dashv G$ respectively, then we say that $\overline{F} \dashv \overline{G}$ is a *lifting of the adjunction* $F \dashv G$ (alternatively, this is an adjunction in $\mathrm{Cat}^{\rightarrow}$, the 2-category of functors and commuting squares [Her93]). ◂

This definition differs from the usual notion of a *fibred adjunction*, as we do not assume fibredness of either adjoint. However, it has been shown in [Win90, Lemma 4.5] for fibrations over a single base category, and later generalised in [Her93, Lemma 3.3.3] to fibrations over arbitrary base categories, that the right adjoint in such a lifting of an adjunction is in fact always fibred. Dually, we have that the left adjoint is co-fibred (i.e., it preserves op-Cartesian maps).

A family of instances is given by predicate and relation liftings.

Lemma 4.3.2: Let $\mathcal{C}, \mathcal{D}$ be complete and well-powered categories with logical factorisation systems $(\mathcal{E}_1, \mathcal{M}_1), (\mathcal{E}_2, \mathcal{M}_2)$. Given an adjunction $F \dashv G \colon \mathcal{D} \to \mathcal{C}$ such that $F(\mathcal{E}_1) \subseteq \mathcal{E}_2$, the predicate and relation liftings form liftings of the adjunction:

(4.6)
$$\begin{tikzcd} \mathsf{Pred}(\mathcal{C}) \arrow[r, bend left, "\mathsf{Pred}(F)"] \arrow[r, phantom, "\perp"] \arrow[d] & \mathsf{Pred}(\mathcal{D}) \arrow[l, bend left, "\mathsf{Pred}(G)"] \arrow[d] \\ \mathcal{C} \arrow[r, bend left, "F"] \arrow[r, phantom, "\perp"] & \mathcal{D} \arrow[l, bend left, "G"] \end{tikzcd} \qquad \begin{tikzcd} \mathsf{Rel}(\mathcal{C}) \arrow[r, bend left, "\mathsf{Rel}(F)"] \arrow[r, phantom, "\perp"] \arrow[d] & \mathsf{Rel}(\mathcal{D}) \arrow[l, bend left, "\mathsf{Rel}(G)"] \arrow[d] \\ \mathcal{C} \arrow[r, bend left, "F"] \arrow[r, phantom, "\perp"] & \mathcal{D} \arrow[l, bend left, "G"] \end{tikzcd}$$ ◀

Proof. See Lemma 3.3.4 for the relation liftings. Lifting to predicates is similar, as relations are predicates on $X \times X$. ■

From these liftings, especially the case for predicates, we can lift also to quotients and equivalence relations given some extra conditions. For a category $\mathcal{C}$, we denote by $\mathrm{Quot}(\mathcal{C})$ the category of co-subobjects of objects of $\mathcal{C}$, that is, equivalence classes of (abstract) epimorphisms. This is exactly the category $\mathsf{Pred}(\mathcal{C}^{\mathrm{op}})^{\mathrm{op}}$, so that also $\mathsf{Pred}(\mathcal{C}^{\mathrm{op}}) \simeq \mathrm{Quot}(\mathcal{C})^{\mathrm{op}}$ and $\mathrm{Quot}(\mathcal{C}^{\mathrm{op}}) \simeq \mathsf{Pred}(\mathcal{C})^{\mathrm{op}}$. Further, we can define quotient lifting dually to predicate lifting. The following is then the dual of the above result.

Corollary 4.3.3: Let $\mathcal{C}, \mathcal{D}$ be co-complete and co-well-powered categories such that $\mathcal{C}^{\mathrm{op}}, \mathcal{D}^{\mathrm{op}}$ have logical factorisation systems $(\mathcal{M}_1, \mathcal{E}_1), (\mathcal{M}_2, \mathcal{E}_2)$. Then, given an adjunction $F \dashv G \colon \mathcal{D} \to \mathcal{C}$, s.t. $G(\mathcal{M}_2) \subseteq \mathcal{M}_1$, the quotient liftings form a lifting of the adjunction:

(4.7)
$$\begin{tikzcd} \mathrm{Quot}(\mathcal{C}) \arrow[r, bend left, "\mathrm{Quot}(F)"] \arrow[r, phantom, "\perp"] \arrow[d] & \mathrm{Quot}(\mathcal{D}) \arrow[l, bend left, "\mathrm{Quot}(G)"] \arrow[d] \\ \mathcal{C} \arrow[r, bend left, "F"] \arrow[r, phantom, "\perp"] & \mathcal{D} \arrow[l, bend left, "G"] \end{tikzcd}$$ ◀

As discussed above, predicates and quotients in the opposite category are (as objects) exactly quotients and predicates in the original category respectively. In the following result, we take a dual adjunction, so that the lifting gives an adjunction between predicates and quotients. We further give conditions under which the quotients correspond to equivalence relations (ERel). We then have adjunctions between predicates and equivalence relations, which we require for our applications to modal logic in Section 4.6.

Corollary 4.3.4: Let $\mathcal{C}$ and $\mathcal{D}$ be complete, well-powered and co-complete, co-well-powered categories respectively, such that $\mathcal{C}$ and $\mathcal{D}^{\mathrm{op}}$ have logical factorisation systems $(\mathcal{E}_1, \mathcal{M}_1)$, $(\mathcal{M}_2, \mathcal{E}_2)$. Suppose also an adjunction $F \dashv G : \mathcal{D}^{\mathrm{op}} \to \mathcal{C}$, such that $F(\mathcal{E}_1) \subseteq \mathcal{M}_2$.

Then the following hold:

- If $\mathcal{D}$ is an exact category (in the sense of Barr) in which all epis are regular, we obtain a lifting of the adjunction as on the left below.
- If instead $\mathcal{C}$ is exact and all epis in $\mathcal{C}$ are regular we obtain the lifting of the adjunction as on the right.

$$\begin{tikzcd} \mathsf{Pred}(\mathcal{C}) \arrow[r, bend left, "\overline{F}"] \arrow[r, phantom, "\perp"] \arrow[d] & \mathsf{ERel}(\mathcal{D})^{\mathrm{op}} \arrow[l, bend left, "\overline{G}"] \arrow[d] \\ \mathcal{C} \arrow[r, bend left, "F"] \arrow[r, phantom, "\perp"] & \mathcal{D}^{\mathrm{op}} \arrow[l, bend left, "G"] \end{tikzcd} \qquad \begin{tikzcd} \mathsf{ERel}(\mathcal{C}) \arrow[r, bend left, "\overline{F}"] \arrow[r, phantom, "\perp"] \arrow[d] & \mathsf{Pred}(\mathcal{D})^{\mathrm{op}} \arrow[l, bend left, "\overline{G}"] \arrow[d] \\ \mathcal{C} \arrow[r, bend left, "F"] \arrow[r, phantom, "\perp"] & \mathcal{D}^{\mathrm{op}} \arrow[l, bend left, "G"] \end{tikzcd} \tag{4.8}$$ ◀

Proof. For the first part, we take predicate fibrations on both sides of the adjunction. The global assumptions mean that the predicate liftings give a lifting of the adjunction. Predicates in $\mathcal{D}^{\mathrm{op}}$ correspond to quotients in $\mathcal{D}$ by definition, and the assumption that $\mathcal{D}$ is exact and all epis of $\mathcal{D}$ are regular means that quotients in turn correspond to equivalence relations.

For the second part, we instead take quotient fibrations on both sides, with the global assumptions ensuring that the quotient liftings give a lifting of the adjunction. Then, as $\mathcal{C}$ is assumed to be exact with all epis in $\mathcal{C}$ regular, the quotients in $\mathcal{C}$ correspond to equivalence relations. Above $\mathcal{D}^{\mathrm{op}}$ we have $\mathsf{Pred}(\mathcal{D})^{\mathrm{op}}$ as predicates are exactly dual to quotients. ■

Our goal is now to relate liftings of adjunctions to adjunctions defined between fibres. In [Jac01] it is shown how this can be done for fibrations over a single base category. Also studied in [Her92; Her93] is how adjunctions between fibrations with distinct base categories arise from adjunctions between fibrations with a common base category.

Lemma 4.3.5: Suppose we have a lifted adjunction as in (4.5). Then we also have the following adjunctions between fibres, for all objects X of $\mathcal{C}$ and

Y of $\mathcal{D}$, where η and ε are the unit and counit of the adjunction $F \dashv G$ respectively.

$$\text{(4.9)} \qquad \begin{tikzcd} \mathcal{E}_X \arrow[r, bend left, "\overline{F}"] \arrow[r, phantom, "\perp"] & \mathcal{F}_{FX} \arrow[l, bend left, "\eta_X^* \circ \overline{G}"] \end{tikzcd} \qquad\qquad \begin{tikzcd} \mathcal{E}_{GY} \arrow[r, bend left, "\coprod_{\varepsilon_Y} \circ \overline{F}"] \arrow[r, phantom, "\perp"] & \mathcal{F}_Y \arrow[l, bend left, "\overline{G}"] \end{tikzcd} \qquad ◂$$

Proof. For the left-hand adjunction, let $x \in \mathcal{E}_X$. Then, we have $x \leq \eta^* \circ \overline{G} \circ \overline{F}(x)$ because there is a map $x \to \overline{G} \circ \overline{F}(x)$ by the assumption of a lifted adjunction and a *Cartesian* map $\eta^* \circ \overline{G} \circ \overline{F}(x) \to \overline{G} \circ \overline{F}(x)$ by definition of η^*.

For any $y \in \mathcal{F}_{FX}$, similarly to above, we have a map $\eta^* \circ \overline{G}(y) \to \overline{G}(y)$ whose transpose under the assumed lifted adjunction yields the inequality $\overline{F} \circ \eta^* \circ \overline{G}(y) \leq y$.

The proof for the right-hand adjunction is similar and a straightforward generalisation of Lemma 3.3.7. ∎

Returning to the example of adjunctions for predicate and relation liftings (Lemma 4.3.2), Lemma 4.3.5 allows us to obtain adjunctions between fibres, which are of interest when we study invariants and coinductive predicates in the coming section. In particular, we recover the adjunctions from Section 4.2.

Example 4.3.6 (Eilenberg-Moore): For the case of an adjunction monadic over Set, each category (Set and $\mathcal{EM}(T)$ for a monad T) has a $(\mathsf{RegEpi}, \mathsf{Mono})$-factorisation system as they are regular. Also, the abstract epis are preserved by left adjoints and these categories are complete and well-powered. We thus obtain a lifting of any monadic adjunction to predicates, relations and quotients by Lemma 4.3.2 and Corollary 4.3.3. Furthermore, the adjunction between fibres as on the right in Eq. (4.9) then exactly instantiates to the adjunctions discussed in Section 4.2.1 for the cases of compact Hausdorff spaces and convex algebras. The left adjoint in each case takes the closure or convex hull of a relation on a set. ◂

In fact, each of these "local" adjunctions implies the existence of the other. Note that we do not assume a (global) lifted adjunction, so we must assume (co-)fibredness explicitly.

Lemma 4.3.7: Suppose we have an adjunction $F \dashv G : \mathcal{D} \to \mathcal{C}$ and bifibrations $p : \mathcal{E} \to \mathcal{C}$, $q : \mathcal{F} \to \mathcal{D}$. Also, suppose $\overline{G}$ is a fibred lifting of G and $\overline{F}$ is a co-fibred lifting of F. Then we have adjunctions $\overline{F}_X \dashv \eta_X^* \circ \overline{G}_{FX}$ for all X iff we have adjunctions $\coprod_{\varepsilon_Y} \circ \overline{F}_{GY} \dashv \overline{G}_Y$ for all Y, that is, the adjunctions in (4.9) can be derived from each other. ◂

Proof. From left to right, let Y be some object of $\mathcal{D}$ and consider the following pasting of adjunctions:

(4.10)
$$\begin{tikzcd} \mathcal{E}_{GY} \arrow[r, bend left, "\overline{F}"] \arrow[r, phantom, "\perp"] & \mathcal{F}_{FGY} \arrow[l, bend left, "\eta^*_{GY} \circ \overline{G}"] \arrow[r, bend left, "\coprod_{\varepsilon_Y}"] \arrow[r, phantom, "\perp"] & \mathcal{F}_Y \arrow[l, bend left, "\varepsilon^*_Y"] \end{tikzcd}$$

The composition of right adjoints is equal to

$$\begin{aligned}
&(\overline{G}\text{ fibred}) && \eta^*_{GY} \circ \overline{G}_{FGY} \circ \varepsilon^*_Y = \eta^*_{GY} \circ (G\varepsilon_Y)^* \circ \overline{G}_Y \\
&(\text{split fib.}) && = (G\varepsilon_Y \circ \eta_{GY})^* \circ \overline{G}_Y \\
&(F \dashv G) && = \overline{G}_Y
\end{aligned}$$

The other direction can be shown dually. ∎

Due to results of [Her93; Her92] on factorisation of fibred adjunctions, we can also go from the existence of adjunctions between fibres (above all objects of our base category) to an adjunction between the total categories. As mentioned before, we drop the requirement of (co)fibredness in the first item, as it follows from the existence of the lifting.

Lemma 4.3.8: Suppose we have an adjunction $F \dashv G : \mathcal{D} \to \mathcal{C}$ and fibrations $p : \mathcal{E} \to \mathcal{C}$ and $q : \mathcal{F} \to \mathcal{D}$. Then the following are equivalent

1. A lifting of the adjunction to $\overline{F} \dashv \overline{G} : \mathcal{F} \to \mathcal{E}$
2. A fibred lifting $\overline{G}$ of G and for each object Y of $\mathcal{D}$, a left adjoint to $\overline{G}_Y : \mathcal{F}_Y \to \mathcal{E}_{GY}$
3. A fibred lifting $\overline{G}$ of G and for each object Y of $\mathcal{D}$, $\overline{G}_Y : \mathcal{F}_Y \to \mathcal{E}_{GY}$ preserves meets. ◂

Proof. 1 $\implies$ 2: Assume a lifting of the adjunction $\overline{F} \dashv \overline{G}$. Then, we have seen that $\overline{G}$ is fibred, and Lemma 4.3.5 gives us the local adjunctions.

2 $\implies$ 1: Now suppose we have a fibred lifting $\overline{G} : \mathcal{F} \to \mathcal{E}$ of G and left adjoints to $\overline{G}_Y$ (which we will call L_Y) for all Y in $\mathcal{D}$. It is shown in [Her93], that we have a lifted adjunction

(4.11)
$$\begin{tikzcd} \mathcal{E} \arrow[r, bend left, "\widetilde{F}"] \arrow[r, phantom, "\perp"] \arrow[d, "p"'] & \mathcal{E} \times_{\mathcal{C}} \mathcal{D} \arrow[l, bend left, "p^*(G)"] \arrow[d, "G^*(p)"] \\ \mathcal{C} \arrow[r, bend left, "F"] \arrow[r, phantom, "\perp"] & \mathcal{D} \arrow[l, bend left, "G"] \end{tikzcd}$$

where $G^*(p)$ is obtained via change of base, i.e., it is the pullback of p along G. Now, by universality we obtain a functor $\hat{G} : \mathcal{F} \to \mathcal{E} \times_{\mathcal{C}} \mathcal{D}$ which sends an object R of $\mathcal{F}$ to $(\overline{G}(R), qR)$.

The assumed left adjoints now also give, for each object Y in $\mathcal{D}$ a left adjoint to $\hat{G}_Y$. We define this by $\hat{F}_Y(R, Y) = L_Y(R)$, which is well-defined because R is above GY by definition of the pullback. In this way, we have obtained adjunctions $\hat{F}_Y \dashv \hat{G}_Y$ for each Y in $\mathcal{D}$. By [Jac01, Lemma 1.8.9] (where we do not require the Beck-Chevalley condition to hold as we do not require fibredness of the left adjoint) we then have a left adjoint $\hat{F} \dashv \hat{G}$. We thus have the following picture:

$$
\begin{tikzcd}
 & & \mathcal{F} \arrow[dll, bend right, "\overline{G}"'] \arrow[dl, bend left, "\hat{G}"] \arrow[ddl, bend left, "q"] \\
\mathcal{E} \arrow[r, bend left, "\widetilde{F}"] \arrow[r, phantom, "\perp"] \arrow[d, "p"'] & \mathcal{E} \times_{\mathcal{C}} \mathcal{D} \arrow[l, bend left, "p^*(G)"] \arrow[ur, bend left, "\hat{F}"] \arrow[ur, phantom, "\perp"] \arrow[d, "G^*(p)"] & \\
\mathcal{C} \arrow[r, bend left, "F"] \arrow[r, phantom, "\perp"] & \mathcal{D} \arrow[l, bend left, "G"] &
\end{tikzcd}
\tag{4.12}
$$

Finally, as shown in [Her93], this gives a left adjoint $\overline{F} = \hat{F} \circ \widetilde{F}$ of $\overline{G}$ by composition of adjoints.

3 $\iff$ 2: As we work in complete lattice fibrations, existence of a left adjoint on the fibres and preservation of meets are equivalent. ∎

4.4 Comparing coinductive predicates along steps

In this section, we consider endofunctors in the setting of an adjunction, and will study coalgebras for these endofunctors—and sometimes algebras, viewed as coalgebras in an opposite category. These endofunctors are connected via the notion of a *step* [RJL21], which is a natural transformation that allows one to transport coalgebras along the adjunction. More formally, steps give rise to liftings of the right and left adjoint (depending on which kind of step) to categories of coalgebras. The key question that we address in this section is whether these liftings to categories of coalgebras preserve a coinductive predicate of interest.

In the following context

$$
\begin{tikzcd}
B \arrow[loop left] & \mathcal{C} \arrow[r, bend left, "F"] \arrow[r, phantom, "\perp"] & \mathcal{D} \arrow[l, bend left, "G"] & L \arrow[from=l, phantom] \arrow[loop right]
\end{tikzcd}
\tag{4.13}
$$

we will again speak of *forward steps* $\delta \colon BG \to GL$ and *backward steps* $\iota \colon GL \to BG$. The induced liftings of the adjoints above will also be denoted $\hat{F} \colon \mathrm{Coalg}(B) \to \mathrm{Coalg}(L)$ and $\hat{G} \colon \mathrm{Coalg}(L) \to \mathrm{Coalg}(B)$.

Remark 4.4.1: The names "forward" and "backward" steps were introduced in the previous chapter. However, there they were assumed to be each other's one-sided inverses. In the current chapter, we make no such assumption and study forward and backward steps independently of each other. We will even consider an example (see below), where one of the transformations may not be natural, meaning we must consider one as just a collection of maps. In [RJL21] only what we refer to as a forward step appears. There is a clear asymmetry between the two; forward steps have a mate correspondence, and at least two other equivalent presentations via transposing along the adjunction. For backward steps there seem to be no such equivalent characterisations, as the left adjoint is on the "wrong" side. ◀

These types of steps occur, e.g., as the one-step semantics of coalgebraic modal logics (more usually, the mate of a forward step) [PMW06; Kli07; JS10; RJL21] The examples of Chapter 3 are still relevant here: those involved in the construction of ultrafilter extensions of coalgebras [KKP05; Gar20], and those combining non-determinism and probabilities [BS21; BSS21; GP20; Goy21].

Example 4.4.2: An example of such transformations not covered in the previous chapter occurs in the presentation of a determinisation procedure for probabilistic automata given in [BSS21]. In the current chapter, we will be able to use a result of *op. cit.*: that there is an injective natural transformation $\iota \colon \mathcal{U} \circ \mathcal{P}_c^\Sigma \to \mathcal{P}^\Sigma \circ \mathcal{U}$, induced by an analogous transformation for $B = \mathcal{P}$ and $L = \mathcal{P}_c$. Such a transformation without labels simply includes convex subsets into the powerset. Injectivity induces componentwise inverses forming, for each subset, its convex hull. These components may not yield a *natural* transformation, requiring our more general approach developed in the sequel. ◀

4.4.1 Comparing coinductive predicates

We have so far seen how steps defined for an adjunction with endofunctors on each of the categories allow us to map coalgebras along this adjunction. Further, when we have fibrations on each of the categories of the adjunction, and liftings of the involved functors, we can define predicates on these coalgebras. Next, we will combine these transformations and give conditions on the steps and liftings, which allow us to link predicates on a coalgebra with predicates on the coalgebra obtained by applying a step-induced lifting.

Assumption 4.4.3: Throughout this section, we assume $\mathsf{CLat}_\sqcap$-fibrations $p\colon \mathcal{E} \to \mathcal{C}$ and $q\colon \mathcal{F} \to \mathcal{D}$ (see Section 2.5.2) in the following setting:

$$\begin{tikzcd} \mathcal{E} \arrow[loop left, "\overline{B}"] \arrow[r, bend left, "\overline{F}"] \arrow[r, phantom, "\perp"] \arrow[d, "p"'] & \mathcal{F} \arrow[l, bend left, "\overline{G}"] \arrow[loop right, "\overline{L}"] \arrow[d, "q"] \\ \mathcal{C} \arrow[loop left, "B"] \arrow[r, bend left, "F"] \arrow[r, phantom, "\perp"] & \mathcal{D} \arrow[l, bend left, "G"] \arrow[loop right, "L"] \end{tikzcd} \tag{4.14}$$

where $\overline{F} \dashv \overline{G}$ is a lifting of the adjunction $F \dashv G$, $\overline{B}$ is a lifting of B, and $\overline{L}$ is a lifting of L. ◀

These assumptions give us coinductive predicates on B-coalgebras, using $\overline{B}$, and on L-coalgebras, using $\overline{L}$. This setting allows us to put conditions on forward and backward steps. These conditions, in turn, allow us to obtain steps at the level of the induced adjunctions between fibres, which puts us back into the setting of Section 4.2. In particular, it allows us to apply Lemma 3.3.2 to preserve the relevant coinductive predicates. We now explain this in more detail for backward and forward steps separately.

Preservation via backward steps

Consider a backward step $\iota\colon\ GL \to BG$. Given an L-coalgebra (Y, ξ) together with this backward step and Assumption 4.4.3, we have the following setting.

$$\begin{tikzcd} \mathcal{E}_{GY} \arrow[loop left, "(\iota_Y \circ G\xi)^* \circ \overline{B}"] \arrow[r, bend left, "\coprod_{\varepsilon_Y} \circ \overline{F}"] \arrow[r, phantom, "\perp"] & \mathcal{F}_Y \arrow[l, bend left, "\overline{G}"] \arrow[loop right, "\xi^* \circ \overline{L}"] \end{tikzcd} \tag{4.15}$$

The greatest fixed point $\nu(\xi^* \circ \overline{L})$ is a coinductive predicate on the L-coalgebra (Y, ξ). The greatest fixed point $\nu((\iota_Y \circ G\xi)^* \circ \overline{B})$ is instead a coinductive predicate on the B-coalgebra obtained by applying the lifting $\hat{G}\colon\ \mathrm{Coalg}(L) \to \mathrm{Coalg}(B)$ induced by the step ι to (Y, ξ). Like in Section 4.2, we ask whether the right adjoint $\overline{G}_Y$ preserves the greatest fixed point, that is, maps $\nu(\xi^* \circ \overline{L})$ to $\nu((\iota_Y \circ G\xi)^* \circ \overline{B})$.

The following result gives a sufficient condition for constructing a step in the above adjunction between fibres; this condition is in terms of the backward step ι and the liftings.

Lemma 4.4.4: For a (backward) step $\iota\colon GL \to BG$ and an L-coalgebra $\xi\colon Y \to LY$:

1. If $\overline{G} \circ \overline{L} \leq \iota^* \circ \overline{B} \circ \overline{G}$, then $\overline{G}_Y \circ \xi^* \circ \overline{L}_Y \leq (\iota_Y \circ G\xi)^* \circ \overline{B}_{GY} \circ \overline{G}_Y$;
2. If $\iota^* \circ \overline{B} \circ \overline{G} \leq \overline{G} \circ \overline{L}$, then $\overline{G}_Y \circ \xi^* \circ \overline{L}_Y \geq (\iota_Y \circ G\xi)^* \circ \overline{B}_{GY} \circ \overline{G}_Y$. ◂

Proof. Given the assumption of Item 1, we have

$$\begin{aligned}
&\text{(Cartesian)} && \overline{G}_Y \circ \xi^* \circ \overline{L}_Y \leq (G\xi)^* \circ \overline{G}_{LY} \circ \overline{L}_Y \\
&\text{(assumption)} && \leq (G\xi)^* \circ \iota_Y^* \circ \overline{B}_{GY} \circ \overline{G}_Y \\
&\text{(split fib.)} && = (\iota_Y \circ G\xi)^* \circ \overline{B}_{GY} \circ \overline{G}_Y
\end{aligned}$$

With the assumptions of Item 2, and using the fibredness of $\overline{G}$, we obtain the opposite inequalities. ■

We note that the condition of Item 1 is equivalent to having a lifting $\overline{\iota}\colon \overline{G}\,\overline{L} \to \overline{B}\,\overline{G}$ of ι, using the existence of a Cartesian lifting of ι. The consequent can be seen as a backward step on the Galois connection of Eq. (4.15), and thus similarly entails a lifting of $\overline{G}$ to post-fixed points (coalgebras) of $\xi^* \circ \overline{L}$. In this sense, $\overline{G}$ preserves invariants and, if it is order-reflecting, will also reflect them (cf. the results of the previous chapter for coalgebraic bisimulations). The inequality of Item 2 often requires further assumptions and more work. We will give some instances where it is satisfied in Section 4.6. Together, the assumptions are equivalent to $\overline{\iota}$ itself being a Cartesian map. See Section 4.5 for a further discussion of these conditions, especially how they relate to the preservation and reflection of *invariants*. For our applications, we are most interested in the following corollary, showing what the above conditions mean for the action of $\overline{G}$ on greatest fixed points.

Corollary 4.4.5: Suppose we have a (backward) step $\iota\colon GL \to BG$. Then for any $\xi\colon Y \to LY$:

1. If $\overline{G} \circ \overline{L} \leq \iota^* \circ \overline{B} \circ \overline{G}$, then $\overline{G}_Y(\nu(\xi^* \circ \overline{L}_Y)) \leq \nu((\iota_Y \circ G\xi)^* \circ \overline{B}_{GY})$;
2. If $\iota^* \circ \overline{B} \circ \overline{G} \leq \overline{G} \circ \overline{L}$, then $\overline{G}_Y(\nu(\xi^* \circ \overline{L}_Y)) \geq \nu((\iota_Y \circ G\xi)^* \circ \overline{B}_{GY})$. ◂

Proof. Apply Lemmata 3.3.2, 4.3.5 and 4.4.4 together with the fact the right adjoint of a lifted adjunction is always fibred. ■

Corollary 4.4.5 thus gives sufficient conditions for $\overline{G}_Y$ to map the greatest fixed point of the coinductive predicate on an L-coalgebra (Y, ξ) to the greatest fixed point of the coinductive predicate on the B-coalgebra $\hat{G}(Y, \xi)$ constructed via ι, in the setting of (4.15).

Remark 4.4.6: It is in fact not necessary that ι is natural; that is, Lemma 4.4.4 and Corollary 4.4.5 go through even if ι is just a collection of arrows. The idea is that we can consider the inequalities $\overline{G}\circ\overline{L} \leq \iota^* \circ \overline{B}\circ\overline{G}$ and $\iota^*\circ\overline{B}\circ\overline{G} \leq \overline{G}\circ\overline{L}$ componentwise. We will apply this idea in Section 4.6.1, where we use the existence of *componentwise* left inverses to ι to obtain the inequalities. ◂

Remark 4.4.7: Note that by the assumed Galois connection, the conclusion of Item 2 in Corollary 4.4.5 is equivalent to

$$\left(\coprod\nolimits_{\varepsilon_Y} \circ \overline{F}\right)(\nu((\iota_Y \circ G\xi)^* \circ \overline{B}_{GY})) \leq \nu(\xi^* \circ \overline{L}_Y) \tag{4.16}$$

so that we may also consider the effect of the left adjoint on coinductive predicates, as we do in Proposition 4.6.3. ◂

Preservation via forward steps

We proceed to focus on forward steps. Recall from Lemma 4.3.5 that the lifted adjunction between fibrations induces two types of adjunctions between fibres; for backward steps we used one of them, for forward steps we use the other. We thus work in the following setting:

$$\begin{tikzcd} \mathcal{E}_X \arrow[loop left, "\gamma^*\circ\overline{B}"] \arrow[r, bend left, "\overline{F}"] \arrow[r, phantom, "\perp"] & \mathcal{F}_{FX} \arrow[l, bend left, "\eta_X^*\circ\overline{G}"] \arrow[loop right, "(\hat{\delta}_X\circ F\gamma)^*\circ\overline{L}"] \end{tikzcd} \tag{4.17}$$

where (X,γ) is a B-coalgebra. We have the following result on constructing steps in this adjunction between fibres.

Lemma 4.4.8: Suppose we have a (forward) step $\delta\colon BG \to GL$, then for $\gamma\colon X \to BX$:

1. If $\delta^*\circ\overline{G}\circ\overline{L} \leq \overline{B}\circ\overline{G}$ and $\overline{B}$ is fibred, then

$$\eta_X^*\circ\overline{G}_{FX}\circ(\hat{\delta}_X\circ F\gamma)^*\circ\overline{L}_{FX} \leq \gamma^*\circ\overline{B}_X\circ\eta_X^*\circ\overline{G}_{FX};$$

2. If $\overline{B}\circ\overline{G} \leq \delta^*\circ\overline{G}\circ\overline{L}$, then

$$\eta_X^*\circ\overline{G}_{FX}\circ(\hat{\delta}_X\circ F\gamma)^*\circ\overline{L}_{FX} \geq \gamma^*\circ\overline{B}_X\circ\eta_X^*\circ\overline{G}_{FX} \qquad ◂$$

Proof. Given the assumptions of Item 1, we have:

$$
\begin{aligned}
& && \eta_X^* \circ \overline{G}_{FX} \circ (\hat{\delta}_X \circ F\gamma)^* \circ \overline{L}_{FX} \\
&\text{(split fib.)} && = \eta_X^* \circ \overline{G}_{FX} \circ (F\gamma)^* \circ \hat{\delta}_X^* \circ \overline{L}_{FX} \\
&\text{(Cartesian)} && \leq \eta_X^* \circ (GF\gamma)^* \circ \overline{G}_{FBX} \circ \hat{\delta}_X^* \circ \overline{L}_{FX} \\
&\text{(split fib.)} && = (GF\gamma \circ \eta_X)^* \circ \overline{G}_{FBX} \circ \hat{\delta}_X^* \circ \overline{L}_{FX} \\
&\text{(nat. } \eta) && = (\eta_{BX} \circ \gamma)^* \circ \overline{G}_{FBX} \circ \hat{\delta}_X^* \circ \overline{L}_{FX} \\
&\text{(split fib.)} && = \gamma^* \circ \eta_{BX}^* \circ \overline{G}_{FBX} \circ \hat{\delta}_X^* \circ \overline{L}_{FX} \\
&\text{(Cartesian)} && \leq \gamma^* \circ \eta_{BX}^* \circ (G\hat{\delta}_X)^* \circ \overline{G}_{LFX} \circ \overline{L}_{FX} \\
&\text{((}*\text{) + split)} && = \gamma^* \circ (B\eta_X)^* \circ \delta_{FX}^* \circ \overline{G}_{LFX} \circ \overline{L}_{FX} \\
&\text{(assumption)} && \leq \gamma^* \circ (B\eta_X)^* \circ \overline{B}_{GFX} \circ \overline{G}_{FX} \\
&(\overline{B} \text{ fibred}) && \leq \gamma^* \circ \overline{B}_X \circ \eta_X^* \circ \overline{G}_{FX}
\end{aligned}
$$

With the assumptions of Item 2, the first two inequalities reverse due to fibredness of $\overline{G}$ (which holds since it is a right adjoint), and the last inequality reverses by definition of reindexing. The one-but-last reverses because of the assumption. In the equality marked by $(*)$ we used a basic property of steps and their mates (Lemma 3.2.8). ■

Similarly to backward steps, we get the following result from Lemmata 3.3.2, 4.3.5 and 4.4.8, giving a sufficient condition for preservation of the coinductive predicate by the right adjoint in (4.17).

Corollary 4.4.9: Suppose we have a (forward) step $\delta : BG \to GL$ and a lifting of the adjunction as in Eq. (4.5). Then for $\gamma : X \to BX$:

1. If $\delta^* \circ \overline{G} \circ \overline{L} \leq \overline{B} \circ \overline{G}$ and $\overline{B}$ is fibred, then

$$\eta_X^* \circ \overline{G}_{FX}(\nu((\hat{\delta}_X \circ F\gamma)^* \circ \overline{L}_{FX})) \leq \nu(\gamma^* \circ \overline{B}_X);$$

2. If $\overline{B} \circ \overline{G} \leq \delta^* \circ \overline{G} \circ \overline{L}$, then

$$\eta_X^* \circ \overline{G}_{FX}(\nu((\hat{\delta}_X \circ F\gamma)^* \circ \overline{L}_{FX})) \geq \nu(\gamma^* \circ \overline{B}_X)$$ ◂

Proof. Apply Lemmata 3.3.2, 4.3.5 and 4.4.8 again using the fibredness of the right adjoint in a lifted adjunction. ■

Remark 4.4.10: Contrary to the case of backward steps (see Remark 4.4.6), for forward steps we use naturality in the proof of Lemma 4.4.8. This is done explicitly for η, and is implicit for δ in the application of Lemma 3.2.8. The above proof is thus markedly more involved than for Lemma 4.4.4, emphasising again the asymmetry between forward and backward steps. ◂

Remark 4.4.11: We have assumed that the adjunction between base categories lifts to the total categories of the fibrations, even though the results in Corollary 4.4.9 and Corollary 4.4.5 are about the adjunctions between fibres. Therefore, one might be tempted to only assume these adjunctions between fibres instead of an adjunction between total categories. However, in Lemma 4.4.4 and Lemma 4.4.8 (on which the aforementioned results rely) we also use that the right adjoint $\overline{G}$ is fibred, and if we additionally assume this then it is equivalent to having an adjunction between the total categories (Lemma 4.3.8). ◂

Remark 4.4.12: Our focus is on the separate analysis of ι and δ. If we assume instead that: ι and δ both exist; their liftings form an isomorphism $\bar{\iota}\colon \overline{G}\,\overline{L} \simeq \overline{B}\,\overline{G}$; and we have a fibred lifting $\overline{G}$ of G such that its restriction to fibres preserves meets, then we have $\overline{G}_Y(\nu(\xi^* \circ \overline{L}_Y)) = \nu((\iota_Y \circ G\xi)^* \circ \overline{B}_{GY})$ where $\iota = p\bar{\iota}$, for any L-coalgebra (Y, ξ). Restricting ourselves to fibrations over a single base category, $B = L$, and a lifting of the identity between the total categories, we recover [Spr+21, Prop. 6.2]. ◂

Remark 4.4.13: We can again rewrite the conclusion of Item 2 in Corollary 4.4.9 to give the inequality:

$$\overline{F}(\nu(\gamma^* \circ \overline{B}_X)) \leq \nu((\hat{\delta}_X \circ F\gamma)^* \circ \overline{L}_{FX}) \tag{4.18}$$

We will briefly return to this inequality as well as the one in Eq. (4.16) when we consider examples. ◂

4.5 Comparing coinductive invariants along steps

So far in this chapter, we have considered almost exclusively comparison of coinductive predicates (greatest fixed points). In this section, we briefly return to the case of coinductive *invariants* (post-fixed points) as we studied in the previous chapter (see also the discussion after Lemma 4.4.4). The following results generalise those of Chapter 3 on preservation and reflection of bisimulations.

Lemma 4.5.1: Given a backward step $\iota\colon GL \to BG$, a coalgebra $\xi\colon Y \to LY$, $R \in \mathcal{F}_Y$ and $S \in \mathcal{E}_{GY}$ we have:

1. if $\overline{G} \circ \overline{L} \leq \iota^* \circ \overline{B} \circ \overline{G}$, then $\overline{G}$ preserves invariants, i.e.,

$$R \leq \xi^* \circ \overline{L}(R) \implies \overline{G}(R) \leq (\iota_Y \circ G\xi)^* \circ \overline{B}(\overline{G}(R));$$

2. if $\iota^* \circ \overline{B} \circ \overline{G} \leq \overline{G} \circ \overline{L}$, then $\coprod_\varepsilon \circ \overline{F}$ preserves invariants:

$$S \leq (\iota_Y \circ G\xi)^* \circ \overline{B}(S) \implies \coprod_\varepsilon \circ \overline{F}(S) \leq \xi^* \circ \overline{L}(\coprod_\varepsilon \circ \overline{F}(S));$$

3. if $\iota^* \circ \overline{B} \circ \overline{G} \leq \overline{G} \circ \overline{L}$ and $\overline{G}$ is order-reflecting, then $\overline{G}$ reflects invariants:

$$\overline{G}(R) \leq (\iota_Y \circ G\xi)^* \circ \overline{B}(\overline{G}(R)) \implies R \leq \xi^* \circ \overline{L}(R).$$ ◂

Proof. 1. Suppose $R \leq \xi^* \circ \overline{L}(R)$, then

$$\overline{G}(R) \leq \overline{G}(\xi^* \circ \overline{L}(R)) \tag{4.19}$$
$$\leq (\iota_Y \circ G\xi)^* \circ \overline{B}(\overline{G}(R)) \tag{4.20}$$

with the first inequality holding by monotonicity of $\overline{G}$, and the second by applying Lemma 4.4.4.

2. The inequality given by Lemma 4.4.4 is equivalent to

$$\coprod_\varepsilon \circ \overline{F} \circ (\iota \circ G\xi)^* \circ \overline{B} \leq \xi^* \circ \overline{L} \circ \coprod_\varepsilon \circ \overline{F}$$

as this is the mate of the inequality given there, induced by the adjunction of Eq. (4.15). Then, using monotonicity, we have

$$\coprod_\varepsilon \circ \overline{F}(S) \leq \coprod_\varepsilon \circ \overline{F} \circ (\iota \circ G\xi)^* \circ \overline{B}(S) \tag{4.21}$$
$$\leq \xi^* \circ \overline{L} \circ \coprod_\varepsilon \circ \overline{F}(S) \tag{4.22}$$

3. We have again by Lemma 4.4.4 that

$$\overline{G}(R) \leq (\iota_Y \circ G\xi)^* \circ \overline{B}(\overline{G}(R)) \tag{4.23}$$
$$\leq \overline{G}(\xi^* \circ \overline{L}(R)) \tag{4.24}$$

The given inequality then follows by order-reflection. ■

The same approach can be taken in the setting of forward steps (recall Lemma 4.4.8), yielding the following.

Lemma 4.5.2: Given a forward step $\delta : BG \to GL$, a coalgebra $\gamma : X \to BX$, $R \in \mathcal{F}_{FX}$ and $S \in \mathcal{E}_X$:

1. if $\delta^* \circ \overline{G} \circ \overline{L} \leq \overline{B} \circ \overline{G}$ and $\overline{B}$ is fibred, then $\eta^* \circ \overline{G}$ preserves invariants:

$$R \leq (\hat{\delta}_X \circ F\gamma)^* \circ \overline{L}(R) \implies \eta^*_X \circ \overline{G}(R) \leq \gamma^* \circ \overline{B}(\eta^* \circ \overline{G}(R)).$$

2. if $\overline{B} \circ \overline{G} \leq \delta^* \circ \overline{G} \circ \overline{L}$, then $\overline{F}$ preserves invariants:

$$S \leq \gamma^* \circ \overline{B}(S) \implies \overline{F}(S) \leq (\hat{\delta} \circ F\gamma)^* \circ \overline{L} \circ \overline{F}(S)$$

3. if $\overline{B} \circ \overline{G} \leq \delta^* \circ \overline{G} \circ \overline{L}$ and $\eta^* \circ \overline{G}$ is order-reflecting, then $\eta^* \circ \overline{G}$ reflects invariants:

$$\eta^*_X \circ \overline{G}(R) \leq \gamma^* \circ \overline{B}(\eta^* \circ \overline{G}(R)) \implies R \leq (\hat{\delta}_X \circ F\gamma)^* \circ \overline{L}(R).$$ ◂

4.6 Examples

4.6.1 Lax liftings

Our first application of the results of the previous section continues on from Section 4.2.1 and Example 4.4.2. In the latter example, we introduced the backward step $\iota \colon \mathcal{U} \circ \mathcal{P}_c^\Sigma \to \mathcal{P}^\Sigma \circ \mathcal{U}$ involved in the second stage of the determinisation of probabilistic automata into belief state transformers. We are now able to apply Corollary 4.4.5 to show the preservation and reflection of bisimilarity by the coalgebra lifting induced by ι.

The construction has its roots in the generalised determinisation procedure of [Sil+13], but requires an alternative approach due to the non-existence of a lifting of the powerset monad to convex algebras. The determinisation starts from a monadic adjunction over Set, and then proceeds in two steps: first a lifting of the left adjoint gives a "determinised" system with algebraic structure; then a lifting of the right adjoint forgets this structure to give a system in Set. Here, we consider the second stage and take a lifting of the adjunction and endofunctors B and L to Rel fibrations as in (4.25), so that we may apply our earlier results to show preservation and reflection of bisimilarity.

$$
\begin{tikzcd}
\mathsf{Rel}(\mathrm{Set}) \arrow[loop left, "\mathsf{Rel}(B)"] \arrow[r, bend left, "\mathsf{Rel}(\mathcal{F})"] \arrow[r, phantom, "\perp"] \arrow[d] & \mathsf{Rel}(\mathcal{EM}(T)) \arrow[loop right, "\mathsf{Rel}(L)"] \arrow[l, bend left, "\mathsf{Rel}(\mathcal{U})"] \arrow[d] \\
\mathrm{Set} \arrow[loop left, "B"] \arrow[r, bend left, "\mathcal{F}"] \arrow[r, phantom, "\perp"] & \mathcal{EM}(T) \arrow[loop right, "L"] \arrow[l, bend left, "\mathcal{U}"]
\end{tikzcd}
\tag{4.25}
$$

The lifting of the right adjoint to coalgebras uses a ι which comes from a so-called *lax lifting* [BSS21]. Given a functor $B \colon \mathrm{Set} \to \mathrm{Set}$, a lax lifting of B is a functor $L \colon \mathcal{EM}(T) \to \mathcal{EM}(T)$ such that there is an injective natural transformation $\iota \colon \mathcal{U} \circ L \to B \circ \mathcal{U}$. This almost brings us back into the setting of Chapter 3, except that we do not quite have an invertible step; only componentwise left inverses can be assumed, with no guarantee that these form a natural transformation. We show the following result for the (slightly different) case of transformations that are componentwise *split mono*, and then show how this is applicable to the example of probabilistic automata.

Lemma 4.6.1: The lifting $\hat{\mathcal{U}} \colon \mathrm{Coalg}(L) \to \mathrm{Coalg}(B)$ induced by a componentwise split mono transformation $\iota \colon \mathcal{U} \circ L \to B \circ \mathcal{U}$ preserves and reflects bisimilarity. ◂

Proof. Recall from Example 4.3.6 that we can lift a monadic adjunction to relations. We need to show the inequalities for ι from Corollary 4.4.5. For this, we use the assumption of componentwise split monicity, as this means we have componentwise left inverses which we call $\delta_X : B\mathcal{U}X \to \mathcal{U}LX$ (note that these components need not form a *natural* transformation).

Now, we have, using the fact that $\mathcal{U}$ preserves abstract epis, that ι lifts to a transformation $\bar{\iota} : \mathsf{Rel}(\mathcal{U}) \circ \mathsf{Rel}(L) \to \mathsf{Rel}(B) \circ \mathsf{Rel}(\mathcal{U})$. Also, the components δ_X lift to components $\overline{\delta}_R : \mathsf{Rel}(B) \circ \mathsf{Rel}(\mathcal{U})(R) \to \mathsf{Rel}(\mathcal{U}) \circ \mathsf{Rel}(L)(R)$ as we can lift the componentwise mates $\hat{\delta}_Y : \mathcal{F}BY \to L\mathcal{F}Y$ as also $\mathcal{F}$ preserves regular epis.

We now obtain the inequality

$$\mathsf{Rel}(\mathcal{U})_{LX} \circ \mathsf{Rel}(L)_X \leq \iota_X^* \circ \mathsf{Rel}(B)_{\mathcal{U}X} \circ \mathsf{Rel}(\mathcal{U})_X \tag{4.26}$$

by definition of reindexing. For the other direction, we have

$$\mathsf{Rel}(B)_{\mathcal{U}X} \circ \mathsf{Rel}(\mathcal{U})_X \leq \delta_X^* \circ \mathsf{Rel}(\mathcal{U})_{LX} \circ \mathsf{Rel}(L)_X \tag{4.27}$$

again by definition of reindexing. Applying ι_X^* to both sides gives

$$\begin{align} \iota_X^* \circ \mathsf{Rel}(B)_{\mathcal{U}X} \circ \mathsf{Rel}(\mathcal{U})_X &\leq \iota_X^* \circ \delta_X^* \circ \mathsf{Rel}(\mathcal{U})_{LX} \circ \mathsf{Rel}(L)_X \tag{4.28} \\ &= \mathsf{Rel}(\mathcal{U})_{LX} \circ \mathsf{Rel}(L)_X \tag{4.29} \end{align}$$

where the equality holds as $\delta_X \circ \iota_X = \mathrm{id}_X$ and $\iota_X^* \circ \delta_X^* = (\delta_X \circ \iota_X)^*$.

As we work with the relation fibration and liftings, the coinductive predicates we preserve and reflect are exactly bisimilarity. ■

Taking $B = \mathcal{P}^\Sigma$ and $L = \mathcal{P}_c^\Sigma$ in the setting of Eq. (4.25) brings us back to the setting of determinisation of probabilistic automata considered in Sections 3.2.2 and 3.4.2 (recall also Example 4.4.2). We can then show that *bisimilarity* is preserved and reflected by the second stage of the construction: forgetting the convex structure.

Proposition 4.6.2: Consider the following adjunction:

$$\begin{tikzcd} \mathsf{Rel}(\mathsf{Set}) \arrow[loop left, "\mathsf{Rel}(\mathcal{P}^\Sigma)"] \arrow[r, bend left, "\mathsf{Rel}(\mathcal{F})"] \arrow[r, phantom, "\perp"] \arrow[d] & \mathsf{Rel}(\mathcal{EM}(\mathcal{D})) \arrow[loop right, "\mathsf{Rel}(\mathcal{P}_c^\Sigma)"] \arrow[l, bend left, "\mathsf{Rel}(\mathcal{U})"] \arrow[d] \\ \mathsf{Set} \arrow[loop left, "\mathcal{P}^\Sigma"] \arrow[r, bend left, "\mathcal{F}"] \arrow[r, phantom, "\perp"] & \mathcal{EM}(\mathcal{D}) \arrow[loop right, "\mathcal{P}_c^\Sigma"] \arrow[l, bend left, "\mathcal{U}"] \end{tikzcd} \tag{4.30}$$

and the transformation $\iota : \mathcal{U}\mathcal{P}_c^\Sigma \to \mathcal{P}^\Sigma\mathcal{U}$, whose components are the inclusion maps. Then for any coalgebra $\xi : Y \to \mathcal{P}_c^\Sigma Y$, we have $\overline{U}_Y(\nu(\gamma^* \circ \overline{L}_Y)) = \nu((\iota_Y \circ U\gamma)^* \circ \overline{\mathcal{P}^\Sigma}_{UY})$. ◂

Proof. This follows from Corollary 4.4.5 by observing that the inclusion maps $\mathcal{U}\mathcal{P}_c^{\Sigma}(\boldsymbol{A}) \to \mathcal{P}^{\Sigma}\mathcal{U}(\boldsymbol{A})$ form a componentwise split mono ι because they are injective, and $\mathcal{U} \circ \mathcal{P}_c^{\Sigma}(\boldsymbol{A})$ is only empty when Σ is empty, in which case also $\mathcal{P}^{\Sigma} \circ \mathcal{U}(\boldsymbol{A})$ is empty. ∎

Recall that in Section 3.4.2, we showed a similar result for behavioural equivalence via the preservation and reflection of bisimilarity in case the functor B preserves weak pullbacks. There, we required the existence of a weak lifting. The definition of the injective ι used above is simpler, but requires the more general results of this section to show preservation and reflection by the coalgebra lifting it induces. We do note that the argument for preservation and reflection of *behavioural equivalence* of Lemma 3.3.16 can still be applied here.

In Section 3.4, we considered the setting of coalgebras for the powerset and Vietoris functors, and the transformation of forgetting topological structure on Vietoris coalgebras. This arose as a step-induced lifting of $U\colon \mathsf{Stone} \to \mathsf{Set}$ analogously to $\mathcal{U}\colon \mathcal{EM}(T) \to \mathsf{Set}$ in the current example. From Lemma 4.5.1 we know that the map $\coprod_{\varepsilon} \circ \overline{\mathcal{F}}$ will preserve invariants. We can see this as a generalisation of Theorem 3.4.3, saying that the topological closure of a bisimulation on $\hat{\mathcal{U}}(\boldsymbol{X}, \gamma)$ is a bisimulation on $(\boldsymbol{X}, \gamma)$, by interpreting the action of $\coprod_{\varepsilon} \circ \overline{\mathcal{F}}$ as closing under general algebraic structure. We can further conclude from the inequality of Eq. (4.16) that the closure of bisimilarity on a $\hat{\mathcal{U}}((X, \alpha), \gamma)$ will be below bisimilarity on $((X, \alpha), \gamma)$. Instantiating again to the setting of convex algebras, we obtain the following.

Proposition 4.6.3: In the setting of Eq. (4.30), for any coalgebra $\xi\colon Y \to \mathcal{P}_c^{\Sigma} Y$ and $\mathcal{P}$-bisimulation R on $\hat{\mathcal{U}}(Y, \xi)$, the convex closure of R is a bisimulation on (Y, ξ). Moreover, the convex closure of bisimilarity on $\hat{\mathcal{U}}(Y, \xi)$ is contained in bisimilarity on (Y, ξ). ◂

4.6.2 Expressivity

In this subsection, we will establish the well-known expressivity of modal logic with respect to $\mathcal{P}$-bisimilarity using our abstract framework. Note that for simplicity we consider (unlabelled) transition systems modelled as coalgebras for the (full) powerset functor $\mathcal{P}\colon \mathsf{Set} \to \mathsf{Set}$.

We relate bisimilarity to the logic defined by the following grammar:

$$\phi ::= \Diamond\left(\bigwedge_{i\in I} \phi_i \wedge \bigwedge_{i\in J} \neg\phi_i\right) \tag{4.31}$$

There are no size restrictions on I and J, so that the collection of formulas forms a proper class. As a consequence, the usual syntax based on initial algebras (living in Set) is not well-founded. Although expressivity of a similar infinitary modal logic was already shown coalgebraically in [Mos99], we include this example as it demonstrates how this fundamental result fits into our general framework. This all leads us to the following contravariant adjoint situation

$$\begin{tikzcd} \mathsf{ERel} \arrow[loop left, "\mathsf{Rel}(\mathcal{P})"] \arrow[r, bend left, "\overline{F}"] \arrow[r, phantom, "\perp"] \arrow[d] & \mathsf{Pred}^{\mathrm{op}} \arrow[l, bend left, "\overline{G}"] \arrow[loop right, "\mathsf{Pred}(L)"] \arrow[d] \\ \mathsf{Set} \arrow[loop left, "\mathcal{P}"] \arrow[r, bend left, "F=2^{-}"] \arrow[r, phantom, "\perp"] & \mathsf{Set}^{\mathrm{op}} \arrow[l, bend left, "G=2^{-}"] \arrow[loop right, "L=\mathcal{P}(2\times-)"] \end{tikzcd} \tag{4.32}$$

Note that in Set, we have the factorisation system (Epi = RegEpi, Mono), which is a logical factorisation system.

Now taking inspiration from [Kli07] where this is done for the finitary case, we consider the coalgebraic modal logic (L, δ), where the syntax is given by the endofunctor $L = \mathcal{P}(2 \times -)$ on Set and the 'one-step' semantics $\delta : \mathcal{P}2^{-} \to 2^{L}$ is defined as follows:

$$\delta_X(S)(U) = \bigvee_{\varphi \in S} \left(\bigwedge_{(1,x)\in U} \varphi(x) \wedge \bigwedge_{(0,x)\in U} \neg\varphi(x) \right) \tag{4.33}$$

Note that the coalgebra lifting of F induced by δ turns a transition system with set X of states into an L-algebra on 2^X. This transformation uses the mate $\hat{\delta} : \mathcal{P}(2 \times 2^{-}) \to 2^{\mathcal{P}}$ (written here as a transformation in Set) of δ, given by the following composition

$$\mathcal{P}(2 \times 2^X) \xrightarrow{\varepsilon_{\mathcal{P}(2\times 2^X)}} 2^{2^{\mathcal{P}(2\times 2^X)}} \xrightarrow{2^{\delta_{2^X}}} 2^{\mathcal{P}\left(2^{2^X}\right)} \xrightarrow{2^{\mathcal{P}\eta_X}} 2^{\mathcal{P}X} \tag{4.34}$$

where $\eta_X(x) = \{U \subseteq X \mid x \in U\}$, and similarly for the counit. After some calculation, we see that this maps a set $U \subseteq 2 \times 2^X$ to the predicate

$$\begin{aligned} & \{T \subseteq X \mid U \in \delta_{2^X}(\{\eta_X(t) \mid t \in T\})\} \\ =& \left\{ T \subseteq X \mid \bigvee_{t\in T} \left(\bigwedge_{(1,S)\in U} S \in \eta_X(t) \wedge \bigwedge_{(0,S)\in U} S \notin \eta_X(t) \right) \right\} && (4.35) \\ =& \{T \subseteq X \mid \exists t \in T. && (4.36) \\ & \quad \forall S \subseteq X. (1, S) \in U \Rightarrow t \in S \wedge (0, S) \in U \Rightarrow t \notin S\} \end{aligned}$$

Thus the step-induced lifting $\hat{F}$ transforms a coalgebra $\gamma\colon X \to \mathcal{P}X$ into the L-algebra $(2^X, \alpha = 2^\gamma \circ \hat{\delta}_X)$ so that

$$\begin{aligned}(4.37)\quad \alpha(U \subseteq 2 \times 2^X) = \{x \in X \mid \exists t \in \gamma(x).\\ \forall S \subseteq X.\,(1,S) \in U \Rightarrow t \in S \wedge (0,S) \in U \Rightarrow t \notin S\}\end{aligned}$$

This shows the sense in which δ provides a one-step semantics; α gives the set of states which satisfy the 'formula' $\Diamond\left(\bigwedge_{(1,S)\in U} S \wedge \bigwedge_{(0,S)\in U} \neg S\right)$ defined by U.

Next, we require the corresponding liftings of the functors in order to invoke Corollary 4.4.9 in proving expressivity of our logic. It is perhaps already intuitively clear what the liftings should do, however, we elaborate here the definitions as they arise from the results of Section 4.3. To obtain the setting of Eq. (4.32), we apply Corollary 4.3.3, whose conditions are met as we work over Set. The involved liftings now come from quotient liftings, but we will describe how they act on the corresponding predicates and relations. We have also chosen to denote the liftings of the adjoints as simply $\overline{F}, \overline{G}$ as they seem to have equal claim to the titles of predicate or quotient lifting. The lifting $\overline{G}$ of $G = 2^-$ takes a predicate $p\colon P \rightarrowtail Y$ in $\mathrm{Pred}(\mathcal{D})^{\mathrm{op}}$ and maps this to the following factorisation

(4.38)
$$\begin{tikzcd} 2^P & & 2^Y \arrow[ll, "2^p"'] \arrow[dl, two heads, "e"] \\ & \overline{G}(P) \arrow[ul, tail, "m"] & \end{tikzcd}$$

Here, the map e sends $\varphi\colon Y \to 2$ to the set

$$(4.39)\qquad \{\psi\colon P \to 2 \mid \forall p \in P.\,\varphi(p) = \psi(p)\}$$

Finally, taking the kernel pair of e gives the desired equivalence relation:

$$(4.40)\qquad \{(\varphi,\psi) \mid \forall p \in P.\,\varphi(p) = \psi(p)\}$$

The lifting which we denote as $\mathrm{Pred}(L)$ in fact also arises as a quotient lifting, but can straightforwardly be seen to act on a predicate as $\mathrm{Pred}(L)(P) = (\mathcal{P}(2 \times P) \rightarrowtail \mathcal{P}(2 \times Y))$.

As (co)inductive predicate, α and $\mathrm{Pred}(L)$ give an abstract notion of definability: precisely those sets $\varphi \in 2^X$ which are "reachable", that is, contained in the least subalgebra in the fibres of all predicates on LX. So we consider the fibrations Pred and ERel of predicates and equivalence relations (see (4.32)) on Set (note ERel is chosen since $\mathcal{P}$-bisimilarity is an equivalence).

In fact, the Galois connection of Section 4.2.2 between the lattices ERel_X and $\mathsf{Pred}_{2X}^{\mathrm{op}}$ can be reconstructed from the adjunction between the total categories ERel and $\mathsf{Pred}^{\mathrm{op}}$. In particular, Lemma 4.3.5 gives: $\overline{F} \dashv \eta_X^* \circ \overline{G} \colon \mathsf{Pred}_{FX}^{\mathrm{op}} \to \mathsf{ERel}_X$. Moreover, $\eta_X^* \circ \overline{G}$ is equal to the upper adjoint g of the aforementioned example, as the following result shows.

Proposition 4.6.4: Let F and G be contravariant powerset functors as above, and $\mathrm{g} \colon \mathsf{Pred}_{FX}^{\mathrm{op}} \to \mathsf{ERel}_X$ as in Eq. (4.3). Then in the adjunction $\overline{F} \dashv \eta_X^* \circ \overline{G} \colon \mathsf{Pred}_{FX}^{\mathrm{op}} \to \mathsf{ERel}_X$ given by Lemma 4.3.5, we have $\eta_X^* \circ \overline{G} = \mathrm{g}$. ◀

Proof. Let $x, y \in X$. Then

$$
\begin{aligned}
&(4.41) && (x,y) \in \eta_X^* \overline{G}(S) &&\iff (\eta_X x, \eta_X y) \in \overline{G}(S)\\
&(4.42) && &&\iff \forall \phi \in S.\, \eta_X(x)(\phi) \leftrightarrow \eta_X(y)(\phi)\\
&(4.43) && &&\iff \forall \phi \in S.\, \phi(x) \leftrightarrow \phi(y)\\
&(4.44) && &&\overset{(4.3)}{\iff} (x,y) \in \mathrm{g}(S).
\end{aligned}
$$

showing the equality by extensionality. ■

Now adequacy and expressivity of our logic L follows by proving their corresponding sufficient condition (cf. Corollary 4.4.9 and the discussion at the end of Section 4.2.2) as in the following proposition.

Proposition 4.6.5: For any $P \rightarrowtail X$, we have

$$\delta^*(\overline{G}\mathsf{Pred}(L)(P)) = \mathsf{Rel}(\mathcal{P})\overline{G}(P)$$ ◀

Proof. First we spell out both sides of the equality, with $P \rightarrowtail X$ and $S, T \subseteq 2^X$:

$$
\begin{aligned}
& && \delta^*(\overline{G}\mathsf{Pred}(L)(P))\\
&(4.45) && = \{(S,T) \mid \forall U \in \mathcal{P}(2 \times P).\, \delta(S)(U) \leftrightarrow \delta(T)(U)\}\\
&(4.46) && = \{(S,T) \mid \forall U \in \mathcal{P}(2 \times P).\\
& && \qquad \left(\begin{array}{c} \bigvee_{\varphi \in S} \left(\bigwedge_{(1,x) \in U} \varphi(x) \wedge \bigwedge_{(0,x) \in U} \neg\varphi(x) \right) \\ \leftrightarrow \bigvee_{\varphi \in T} \left(\bigwedge_{(1,x) \in U} \varphi(x) \wedge \bigwedge_{(0,x) \in U} \neg\varphi(x) \right) \end{array} \right) \Bigg\},
\end{aligned}
$$

and

$$
(4.47) \quad \mathsf{Rel}(\mathcal{P})\overline{G}(P) = \left\{ (S,T) \mid \left(\begin{array}{c} \forall \varphi \in S.\, \exists \psi \in T.\, \forall x \in P.\, \varphi(x) \leftrightarrow \psi(x) \\ \wedge\, \forall \varphi \in T.\, \exists \psi \in S.\, \forall x \in P.\, \varphi(x) \leftrightarrow \psi(x) \end{array} \right) \right\}.
$$

To show the inequality from left to right (used to show expressivity), let (S, T) be a pair in $\delta^*(\overline{G}\mathsf{Pred}(L)(P))$ and suppose, towards a contradiction, that $(S, T) \notin \mathsf{Rel}(\mathcal{P})\overline{G}(P)$. This means there exists either $\varphi \in S$ or $\varphi \in T$ violating the associated condition; assume the first (the other case is the same). Thus, suppose there is $\varphi \in S$ such that for all $\psi \in T$ either there is an element $x_\varphi \in P$ with $\varphi(x) \wedge \neg\psi(x)$ or $x_\varphi \in P$ with $\neg\varphi(x) \wedge \psi(x)$. Choosing such an element of P for each ψ, we obtain a set

$$U = \{(1, x_\varphi) \mid \varphi(x)\} \cup \{(0, x_\varphi) \mid \neg\varphi(x)\} \tag{4.48}$$

Now for this set U we have

$$\bigwedge_{(1,x)\in U} \varphi(x) \wedge \bigwedge_{(0,x)\in U} \neg\varphi(x) \tag{4.49}$$

However, the formula

$$\bigvee_{\psi\in T}\left(\bigwedge_{(1,x)\in U} \psi(x) \wedge \bigwedge_{(0,x)\in U} \neg\psi(x)\right) \tag{4.50}$$

evaluates to false, by construction of U. Therefore, we have

$$(S, T) \notin \delta^*(\overline{G}\mathsf{Pred}(L)(P)) \tag{4.51}$$

which is a contradiction with our assumption.

To show the inequality from right to left (used to show adequacy), let (S, T) be a pair in $\mathsf{Rel}(\mathcal{P})\overline{G}(P)$. We will show that $U \in \delta(T)$ whenever $U \in \delta(S)$, for any $U \subseteq 2 \times P$; the symmetric case is similar. Since $U \in \delta(S)$ we find some $\varphi \in S$ such that

$$\bigwedge_{(1,x)\in U} \varphi(x) \wedge \bigwedge_{(0,x)\in U} \neg\varphi(x) \tag{4.52}$$

holds. Moreover, by relation lifting of $\mathcal{P}$, we also find some $\psi \in T$ equivalent to φ. I.e., the formula

$$\bigwedge_{(1,x)\in U} \psi(x) \wedge \bigwedge_{(0,x)\in U} \neg\psi(x) \tag{4.53}$$

holds. Thus, $U \in \delta(T)$. ■

Corollary 4.6.6: The logic defined by $\delta\colon \mathcal{P}2^- \to 2^L$ as in Eq. (4.33) is adequate and expressive with respect to bisimilarity on $\mathcal{P}$-coalgebras. ◂

Proof. Proposition 4.6.5 shows that the map g of Section 4.2.2, maps the definable predicates of the logic to bisimilarity. As explained in the example, the image of the definable predicates under g relates those states satisfying the same formulas of the logic. This is precisely logical equivalence, which we thus show to be equal to bisimilarity. Equality of these relations is exactly the combination of adequacy and expressivity. ∎

Left adjoints and invariants We conclude this example with a brief discussion of our observations in Remark 4.4.13 and Lemma 4.5.2. The first is the transposed inequality obtained by applying the observation of Remark 4.4.13. Here, this entails that

$$\overline{F}(\nu(\gamma^* \circ \mathsf{Rel}(\mathcal{P})_X)) \supseteq \nu((\hat{\delta}_X \circ 2^\gamma)^* \circ \mathsf{Pred}(L)_{2^X}). \tag{4.54}$$

Note the reversed inclusion as we are in $\mathsf{Pred}^{\mathrm{op}}$. Recalling the action of $\overline{F}$ from Section 4.2.2, we see that the left-hand side consists of those predicates which are closed under bisimilarity. The right-hand side is simply the definable predicates. We therefore see that this inequality says that definable predicates are closed under bisimilarity, recovering this known equivalent characterisation of adequacy of the logic.

Another property related to adequacy arises from Lemma 4.5.2, and the inequality $\mathsf{Rel}(\mathcal{P})\overline{G}(P) \leq \delta^*(\overline{G}\mathsf{Pred}(L)(P))$. From these facts, we can conclude that if R is a post-fixed point of $\gamma^* \circ \mathsf{Rel}(\mathcal{P})$ (a bisimulation), then in the lattice $\mathsf{Pred}_{FX}^{\mathrm{op}}$ we have

$$\overline{F}(R) \leq (\hat{\delta}_X \circ F\gamma)^* \circ \mathsf{Pred}(L)_{FX}(\overline{F}(R)) \tag{4.55}$$

Dualising, this means that in the lattice Pred_{FX}:

$$\overline{F}(R) \supseteq \coprod\nolimits_{2^\gamma \circ \hat{\delta}_X}(\mathcal{P}(2 \times \overline{F}(R))) \tag{4.56}$$

The predicate $\overline{F}(R)$ consists of those subsets of X which are R-closed. The above inclusion thus means that, for any $U \subseteq 2 \times \overline{F}(R)$, i.e., $U = \{(b, S) \mid b \in \{0,1\}, S \subseteq X \text{ closed under } R\}$, the set of states satisfying $\Diamond(\bigwedge_{(1,S)\in U} S \wedge \bigwedge_{(0,S)\in U} \neg S)$ is again closed under the bisimulation R.

4.6.3 Apartness

In this subsection, we again consider (unlabelled) transition systems and rather show how our framework allows us to prove the dual of the Hennessy-Milner theorem: *two states are $\mathcal{P}$-apart [GJ21; Geu22] (classically equivalent to being non-bisimilar) iff there is a formula distinguishing them.*

The notion of $\mathcal{P}$-apartness will be an *apartness relation*. This is a relation R on a set X which is irreflexive, symmetric, and co-transitive (i.e., $\forall x, y \in X.\, x \ R \ y \to \forall z \in Z.\,(x \ R \ z \vee y \ R \ z)$). Following [GJ21], the fibration of apartness relations on Set can be seen as the fibred opposite of ERel, essentially the fibration whose fibres are the lattices ERel_X with their order reversed (see [Jac01] for the details of the construction). In particular, the functor $\neg$ maps a tuple (X, R) (when R is an equivalence/apartness on X) to the tuple $(X, \neg R)$. The setting of interest thus becomes the following:

$$\begin{tikzcd}
\mathsf{ERel}^{\mathsf{fop}} \arrow[r, bend left, "\neg"] \arrow[d] & \mathsf{ERel} \arrow[l, bend left, "\neg"] \arrow[r, bend left, "\overline{F}"] \arrow[r, phantom, "\perp"] \arrow[d] & \mathsf{Pred}^{\mathsf{op}} \arrow[l, bend left, "\overline{G}"] \arrow[d] \arrow[loop right, "\mathsf{Pred}(L)"] \\
\mathsf{Set} \arrow[loop left, "\mathcal{P}"] \arrow[r, equal] & \mathsf{Set} \arrow[r, bend left, "F=2^-"] \arrow[r, phantom, "\perp"] & \mathsf{Set}^{\mathsf{op}} \arrow[l, bend left, "G=2^-"] \arrow[loop right, "L=\mathcal{P}(2\times -)"]
\end{tikzcd} \tag{4.57}$$

Note that the right-hand part of the above diagram is the same as in the previous example so that the action of $\neg\overline{G}$ on a predicate $P \rightarrowtail X$ is simply given by applying negation to Eq. (4.40), yielding

$$\neg\overline{G}(P) = \{(\varphi, \psi) \mid \exists p \in P.\, \varphi \#_p \psi\} \tag{4.58}$$

where $\varphi \#_p \psi \iff [\varphi(p) \wedge \neg\psi(p)] \vee [\neg\varphi(p) \wedge \psi(p)]$.

For the lifting $\mathsf{ERel}^{\mathsf{fop}}(\mathcal{P})$, we consider the following definition[1] which can also be obtained by applying negation to the usual relation lifting of $\mathcal{P}$:

$$\begin{aligned} U \ \mathsf{ERel}^{\mathsf{fop}}(\mathcal{P})(R) \ V \iff & [\exists x \in U.\, \forall y \in V.\, x \ R \ y] \vee \\ & [\exists y \in V.\, \forall x \in U.\, x \ R \ y] \end{aligned} \tag{4.59}$$

Now we are in the position to use Corollary 4.4.9 and establish the dual Hennessy-Milner theorem, which was also recently shown in [Geu22] for image-finite transition systems.

Proposition 4.6.7: For any $P \rightarrowtail X$

$$\mathsf{ERel}^{\mathsf{fop}}(\mathcal{P})(\overline{G}(P)) = \delta^*(\neg\overline{G}(\mathsf{Pred}(L)(P))) \tag{4.60}$$ ◂

Proof. First we compute $\mathsf{ERel}^{\mathsf{fop}}(\mathcal{P})(\overline{G}(P))$ as follows:

$$\mathsf{ERel}^{\mathsf{fop}}(\mathcal{P})(\overline{G}(P)) = \left\{(S, T) \mid \begin{pmatrix} \exists \varphi \in S.\, \forall \psi \in T.\, \exists p \in P.\,(\varphi \#_p \psi) \\ \vee \ \exists \varphi \in T.\, \forall \psi \in S.\, \exists p \in P.\,(\varphi \#_p \psi) \end{pmatrix}\right\}. \tag{4.61}$$

[1] Note that our definition differs from the lifting of an apartness relation given in [GJ21], where the two logical formulae ($\exists x \in U.\, \forall y \in V.\, x \ R \ y$ and $\exists y \in V.\, \forall x \in U.\, x \ R \ y$) are composed incorrectly by conjunction.

We start with the dual of expressivity, i.e., that if two states are $\mathcal{P}$-apart, then there is a distinguishing formula between them. This will follow from Corollary 4.4.9(1), so that we need to show that

$$\delta^*(\neg\overline{G}\mathsf{Pred}(L)(P)) \supseteq \mathsf{ERel}^{\mathsf{fop}}(\mathcal{P})(\overline{G}(P)).$$

Notice the reversal of order since we are working in the lattice $\mathsf{ERel}^{\mathsf{op}}_{\mathcal{P}(2^X)}$.

So suppose $(S,T) \in \mathsf{ERel}^{\mathsf{fop}}(\mathcal{P})(\overline{G}(P))$. Then, we have some $\varphi \in S$ such that $\forall \psi \in T.\, \exists x \in P.\, (\varphi \#_x \psi \vee \psi \#_x \varphi)$ (the proof of the symmetric case is similar). Consider the set

$$U = \{(1,x) \mid \varphi(x)\} \cup \{(0,x) \mid \neg\varphi(x)\} \tag{4.62}$$

Clearly, $U \in \delta(S)$. Furthermore, either $\varphi \#_x \psi$ or $\psi \#_x \varphi$ holds, for some $x \in P$. Thus, $\neg\psi(x)$ ($\psi(x)$) holds in the former (resp. latter) case. In other words,

$$\bigwedge_{\psi \in T}\Big(\bigvee_{(1,x)\in U} \neg\psi(x) \vee \bigvee_{(0,x)} \psi(x)\Big) = 1 \implies U \notin \delta(T). \tag{4.63}$$

For the dual of adequacy, we will use Corollary 4.4.9(2), requiring the inclusion

$$\delta^*(\neg\overline{G}\mathsf{Pred}(L)(P)) \subseteq \mathsf{ERel}^{\mathsf{fop}}(\mathcal{P})(\overline{G}(P))$$

To show this, first let $S, T \subseteq 2^X$ and $(\delta(S), \delta(T)) \in \mathsf{ERel}^{\mathsf{fop}}(\mathcal{P})(\overline{G}(P))$. I.e., there is a $U \subseteq 2 \times X$ such that $(\delta(S) \#_U \delta(T)) \vee (\delta(T) \#_U \delta(S))$. Assume the former case holds (the proof when the latter holds is similar). Thus, $U \in \delta(S)$ and $U \notin \delta(T)$. I.e., we find some $\varphi \in S$ such that $(\bigvee_{(1,x)\in U} \neg\psi(x) \vee \bigvee_{(0,x)\in U} \psi(x))$ holds, for every $\psi \in T$. Clearly, $\neg\psi(x) \implies \varphi(x)$ and $\psi(x) \implies \neg\varphi(x)$. Thus, we find $\varphi \#_x \psi$ or $\psi \#_x \varphi$. ■

Corollary 4.6.8: Two states of a $\mathcal{P}$-coalgebra are $\mathcal{P}$-apart iff there is a formula in the logic defined by $\delta : \mathcal{P}2^{-} \to 2^L$ as in Eq. (4.33) distinguishing them. ◂

Proof. The equality of Proposition 4.6.7 shows that the set of definable predicates of the logic is mapped by ¬g to $\mathcal{P}$-apartness. The image of this map is given in Eq. (4.58) so that the equality implied by Proposition 4.6.7 indeed shows the correspondence between apartness and the existence of distinguishing formulas. ■

4.7 Conclusion

In the previous chapter, we studied a preservation result assuming both a forward step δ and a backward step ι, which form one-sided inverses, that is, $\delta \circ \iota = \mathrm{id}$. In the current chapter, we have considered preservation of coinductive predicates by forward and backward steps as separate cases, which we realise by formulating the conditions in a purely fibrational way instead of assuming inverses. This allows us, for instance, to provide a general preservation result for lax liftings (Section 4.6.1). This could not be obtained from the results of the previous chapter: there we required a *natural* inverse δ, which is not part of the assumptions of a lax lifting (only componentwise inverses are assumed). A natural inverse can in fact be non-trivial to provide; for instance, we may require weak distributive laws for their existence. By moving from the relation fibration to general $\mathsf{CLat}_\sqcap$-fibrations, our results are applicable to fibrations of predicates and equivalence relations, providing a new approach for showing expressivity of coalgebraic modal logics.

In cases where one does have an invertible step, such as in the setting of (weak) liftings, and the predicates of interest are given by canonical relation lifting, the conditions of the previous chapter are still likely preferable. Preservation and reflection of predicates by right adjoints then requires the verification only of preservation of abstract epis in the base categories, instead of the explicit proof of the inequalities of this chapter.

In [KR21] a general approach to expressivity of logics with respect to coinductive predicates is proposed. In that paper, there is a fibration only on one of the two categories, and the coinductive predicate of interest is related to logical equivalence. Logical equivalence is defined there via the semantics of the logic, which is in turn obtained via the universal property of an initial algebra. In contrast, in the current chapter, we have not used initial algebras and instead obtain logical equivalence by characterising "modally definable" on the coalgebra of interest, which yields a notion of logical equivalence by applying the right adjoint in a Galois connection between equivalence relations and predicates. This Galois connection was also used in [Kli05], and in the recent [Beo+23] as the starting point for proving expressivity. Here, instead, we obtain this Galois connection from an adjunction between fibrations.

Future Work In [Beo+23] it was shown how the functional characterising bisimilarity can be synthesised from a 'logic' function. Using the notation of this paper, this means defining $\overline{B}$ in terms of $\overline{L}$. Such constructions and, symmetrically, constructing $\overline{L}$ from $\overline{B}$, are of interest at the global level of contravariant adjunctions. In the context of modal logic, providing such constructions, and understanding when they provide expressive logics may allow conditions for expressivity (or the existence of expressive logics) based only on the functors defining the systems or logics of interest (cf. the condition of abstract epi preservation in Chapter 3).

Last, it would be interesting to try to apply our results on comparing coinductive predicates and lifting adjunctions to the quantitative setting of (pseudo-)metrics. We would then also like to understand the relation between our conditions for expressivity of logics and those based on lax extensions, Kantorovich functors and notions of density in [WS20; Gon+23; For+23]. For instance, do the types of predicate liftings identified in those works induce steps satisfying our conditions?

Part II

Constructing Evidence of Different Behaviours

CHAPTER 5

Proving Behavioural Apartness

Difference in a bit of monstrous
Loch? (9)

5.1 Introduction

As we have seen in the preceding chapters, various notions of equivalence have been defined on coalgebras in their development as a general approach to state-based systems. More recently, Geuvers & Jacobs [GJ21] have investigated notions of inequivalence (or *distinguishability*), which they call apartness. This is done first on concrete examples, such as labelled transition systems, for which definitions of apartness are given and shown to be dual to existing definitions of bisimilarity. Further, the definitions are shown to give rise to inductive proof rules for deriving differences in behaviours. They later show how the definitions of apartness and corresponding proof rules can be obtained generally.

The starting point for this approach is coalgebraic bisimilarity defined using canonical relation lifting. As is done for concrete systems, relations are essentially defined to be an apartness if their complement is a bisimulation. To define this using a lifting, the *opposite* relation lifting is used, which takes the complement of a relation, then applies the usual relation lifting, and then complements again. The associated inductive predicate recovers the concrete definitions. It is also shown how proof rules can be obtained recursively for a class of functors generated from a given grammar.

The definitions and proof rules in all these instances are such that states can be distinguished by providing *finite* evidence for a difference in their behaviours, in keeping with the constructive origins of apartness. This is in contrast to the situation for proof systems of coinductive notions, due to their inherent circularity [Luc+09; Clo+16; Bas18; SM17]. However, the functors covered by [GJ21] do not include one capturing probabilistic behaviours, such as the subdistribution functor. It in fact turns out that for this functor, the definition based on opposite relation lifting does not (immediately) yield a satisfactory proof system, in the sense that finite evidence suffices to show apartness. We discuss this issue in Section 5.2: showing apartness of states in a coalgebra for the subdistribution functor requires one to show that *all* couplings between successor distributions satisfy a certain property. There are infinitely many such couplings in general, and it is not evident how to simplify the situation directly.

Instead, we turn to an alternative existing notion of equivalence on coalgebras: behavioural equivalence, which relates states that can be identified under a coalgebra homomorphism. Using a characterisation of behavioural equivalence due to Aczel and Mendler based on so-called *precongruences* [AM89], we obtain a new definition of apartness in probabilistic transition systems. For coalgebras for the subdistribution functor, the definition is dual to the well-known probabilistic bisimilarity of Larsen and Skou [LS91], which has been studied coalgebraically already in [VR99; Sok05; Sok11]. This definition will allow us to give finite witnesses in establishing apartness of states.

We are then able to provide a sound and complete derivation system for these notions of apartness on coalgebras for finitary functors. Completeness with respect to apartness on a given coalgebra will be shown by relating proofs in our system to finite-step behaviours, as captured by the final chain of the coalgebra's functor, say B. This chain is defined by transfinite induction. Its objects indexed by finite ordinals are finite applications of B to the one-element set 1. An object $B^n 1$ can then be thought of as containing all possible n-step behaviours of states in a B-coalgebra, and it is known that a coalgebra $\gamma : X \to BX$ induces maps $f_n : X \to B^n 1$, mapping a state to its n-step behaviour. The essence of our completeness proof is then to show that states which are distinguished by a map f_n can be proved apart by a derivation in our system of depth n. In this way, the restriction to finitary functors also becomes clear: apartness proofs (which are inherently finite depth) can be given for those systems whose behaviour is captured by the finite stages $B^n 1$ of the final chain.

Outline and Contributions

We first recall the approach to defining apartness on coalgebras of [GJ21], and exhibit the difficulty encountered when instantiated to probabilistic systems (Section 5.2).

Our first main result is then a general definition of apartness in coalgebras, obtained by dualising a definition of coalgebraic behavioural equivalence. We then provide corresponding proof rules, and show them to be sound and complete with respect to the definition for coalgebras of finitary functors. Instantiating the rules to coalgebras for the subdistribution functor, we show that they do not share the problem encountered when using opposite relation lifting (Section 5.3).

We go on to improve the usability of the proof rules by showing that a proof step showing apartness of a pair of states need only consider their immediate successors. Finally, we give an inductive characterisation of the property which needs to be shown of these successors, so that proof rules can be obtained recursively for functors built from a given grammar (Section 5.4).

All of our results are also illustrated by examples, including a proof of apartness of states in a stream system with an infinite state space, showing the importance of our results in Section 5.4.

Notation

For an equivalence relation $R \subseteq X \times X$, we denote the quotient map by $q_R : X \to X/R$, which sends an element x to its equivalence class $[x]_R$. We extend this to arbitrary relations $R \subseteq X \times X$, by defining q_R as $q_{e(R)}$, where $e(R)$ is the equivalence closure of R. Further, we write $\overline{R}$ for the complement of a relation $R \subseteq X \times X$ on a set X.

5.2 Cobisimilarity

We start by recalling the well-known notion of coalgebraic bisimulation, defined via relation lifting [HJ98]. We have already used this form of bisimulation, but present it here in a proof rule-like form similarly to [GJ21], where the dual notion is developed using the so called fibred opposite construction. In this section, we review this approach and show how it can be utilised to obtain a proof system for apartness. We will further show that there exist functors for which this dualised definition does not give rise to a satisfactory proof system, in contrast to the rules provided in *op. cit.*

Note that we will use alternative terminology to *op. cit.* *Apartness* will be used to mean the dual of equivalence, both in the general sense, and in the formal dual of equivalence relations (in [GJ21] these are referred to as *proper* apartness relations). For dual notions to specific equivalences, we use new terminology. Cobisimulation and cobisimilarity, introduced in this section, will be the duals of (coalgebraic) bisimulation and bisimilarity. Later, we call the dual of coalgebraic behavioural equivalence, "behavioural apartness".

Assuming an endofunctor B on Set, and its canonical relation lifting $\mathsf{Rel}(B)$, recall the following notion of equivalence:

Definition 5.2.1: A relation $R \subseteq X \times X$ is a coalgebraic bisimulation for a coalgebra $\gamma\colon X \to B(X)$ if $R \subseteq (\gamma \times \gamma)^{-1} \circ \mathsf{Rel}(B)(R)$, that is, it satisfies:

$$\frac{x\ R\ y}{\gamma(x)\ \mathsf{Rel}(B)(R)\ \gamma(y)} \tag{5.1}$$

The largest such relation is called *bisimilarity*, denoted by $\underline{\leftrightarrow}$. ◂

Replacing R above with its complement and taking the contrapositive, we obtain a dual definition.

Definition 5.2.2: A *cobisimulation* on a coalgebra $\gamma\colon X \to B(X)$ is a relation R satisfying

$$\frac{\gamma(x)\ \overline{\mathsf{Rel}(B)(\overline{R})}\ \gamma(y)}{x\ R\ y} \tag{5.2}$$

Equivalently, R is a cobisimulation if $R \supseteq (\gamma \times \gamma)^{-1} \circ \overline{\mathsf{Rel}(B)(\overline{R})}$, where the right-hand relation is constructed as the following composition of monotone maps:

$$\mathsf{Rel}_X \xrightarrow{\overline{(-)}} \mathsf{Rel}_X \xrightarrow{\mathsf{Rel}(B)} \mathsf{Rel}_{B(X)} \xrightarrow{\overline{(-)}} \mathsf{Rel}_{B(X)} \xrightarrow{(\gamma\times\gamma)^{-1}} \mathsf{Rel}_X \tag{5.3}$$

in which Rel_X denotes the lattice of relations on X, ordered by inclusion. ◂

For a fibrational treatment of this mode of definition using the *fibred opposite* construction, see [GJ21].

Dually to the familiar case of bisimilarity, the relation of *cobisimilarity* is then defined as the smallest cobisimulation.

Definition 5.2.3: For a coalgebra $\gamma\colon X \to BX$, *cobisimilarity* is the smallest relation $R \subseteq X \times X$ which is a cobisimulation on γ. We denote this by $\underline{\#}$. ◂

By this definition, we see that $\underline{\leftrightarrow} = \overline{(\underline{\#})}$, i.e., bisimilarity and cobisimilarity are each other's complement. Following from this duality is also the fact that, where bisimilarity is an *equivalence* relation, cobisimilarity is an *apartness* relation.

Definition 5.2.4: A relation $R \subseteq X \times X$ on a set X is an *apartness* relation if the following conditions hold:

1. $\forall x \in X.\, \neg(x\ R\ x)$
2. $\forall x, y \in X.\, x\ R\ y \implies y\ R\ x$
3. $\forall x, y, z \in X.\, x\ R\ y \implies x\ R\ z \vee y\ R\ z$ ◂

These conditions will be called *irreflexivity*, *symmetry*, and *cotransitivity* respectively. Clearly, if R is an apartness relation, then its complement $\overline{R}$ will satisfy the dual axioms of reflexivity, symmetry and transitivity, and will thus be an equivalence relation.

In later sections it will be important that we are able to construct, from an arbitrary relation R, a canonical apartness relation. This can be done similarly to the well-known equivalence closure.

Definition 5.2.5: For a relation $R \subseteq X \times X$, the *equivalence closure* of R (denoted by $e(R)$), is the smallest equivalence relation containing R. Dually, the *apartness interior* of R (denoted by R°) is the largest apartness relation contained in R. ◂

These constructions interact in the expected way with complement. For instance, we have $e(\overline{R}) = \overline{R^\circ}$.

Example 5.2.6 (Labelled Transition Systems): In [GJ21] cobisimilarity dual to weak forms of bisimulation are studied. To simplify the presentation, we instantiate the definition of cobisimulation to LTSs, which we model as coalgebras for the functor $\mathcal{P}(-)^\Sigma$ for an input alphabet Σ.

The relation lifting of $\mathcal{P}(-)^\Sigma$ acts as follows

$$\begin{aligned} &\mathsf{Rel}(\mathcal{P}(-)^\Sigma)(R) \\ &= \{(f, g) \mid \forall a \in \Sigma.[\forall x \in f(a).\, \exists y \in g(a).\, (x, y) \in R] \wedge \\ &\qquad [\forall y \in g(a).\, \exists x \in f(a).\, (x, y) \in R]\} \end{aligned} \tag{5.4}$$

Applying the dualisation as in (5.2), we see that a relation R is a cobisimulation on a coalgebra $\gamma : X \to \mathcal{P}(X)^\Sigma$ if it satisfies the following two rules

$$\frac{x \xrightarrow{a} x' \qquad \forall y'.\, y \xrightarrow{a} y' \implies x'\ R\ y'}{x\ R\ y} \tag{5.5}$$

$$\frac{y \xrightarrow{a} y' \qquad \forall x'.\, x \xrightarrow{a} x' \implies x'\ R\ y'}{x\ R\ y} \tag{5.6}$$

where $x \xrightarrow{a} x'$ means $x' \in \gamma(x)(a)$. Thus, apartness of states x, y can be shown by giving a successor of x which is apart from all successors of y. Then cobisimilarity on LTSs is the smallest relation satisfying these rules. This gives an inductive proof system for cobisimilarity, as further explained in [GJ21].

We finish this example by giving a derivation of cobisimilarity for the states x and y in the following LTS:

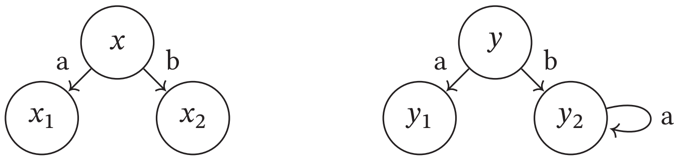

This derivation goes as follows:

$$\frac{x \xrightarrow{b} x_2 \qquad \dfrac{\dfrac{y_2 \xrightarrow{a} y_2 \qquad \dfrac{}{\forall x'.\, x_2 \xrightarrow{a} x'.\, y_2 \# x'}}{x_2 \# y_2}}{\forall y'.\, y \xrightarrow{b} y'.\, x_2 \# y'}}{x \# y} \tag{5.7}$$

The states x and y can be distinguished by their b-transitions, and their successors by the presence of an a-transition only on y_2. ◂

Example 5.2.7 (Subdistributions): We continue by instantiating rule (5.2) for cobisimulations to coalgebras for the subdistribution functor. In contrast to LTSs above, the relation lifting approach does not give a pleasant proof system.

The (finitely supported) subdistribution functor $\mathcal{D}_s :$ Set $\rightarrow$ Set is defined as follows, where $\operatorname{supp}(\mu) = \{x \in X \mid \mu(x) \neq 0\}$.

$$\mathcal{D}_s(X) = \left\{ \mu : X \rightarrow [0,1] \mid \sum_{x \in X} \mu(x) \leq 1, \operatorname{supp}(\mu) \text{ finite} \right\} \tag{5.8}$$

We may equivalently write such distributions as formal sums $\sum_{x \in X} \mu(x) \,|x\rangle$. This allows us to denote the functor's action on arrows as:

$$\mathcal{D}_s(f : X \rightarrow Y) : \left(\sum_{x \in X} \mu(x) \,|x\rangle \right) \mapsto \left(\sum_{x \in X} \mu(x) \,|f(x)\rangle \right) \tag{5.9}$$

We can now elaborate the action of the relation lifting of the subdistribution functor $\mathcal{D}_s$ on a relation $R \subseteq X \times X$:

$$\begin{aligned}
&(5.10) && \mathsf{Rel}(\mathcal{D}_s)(R) \\
&(5.11) && = \{(\mathcal{D}_s(\pi_1)(z), \mathcal{D}_s(\pi_2)(z)) \mid z \in \mathcal{D}_s(R)\} \\
&(5.12) && = \left\{ \left[\sum_{x \in X} \left[\sum_{y \in X} \mu(x,y) \right] |x\rangle \,, \sum_{y \in X} \left[\sum_{x \in X} \mu(x,y) \right] |y\rangle \right] \mid \mu \in \mathcal{D}_s(R) \right\}
\end{aligned}$$

Here, μ is a subdistribution over the elements of R, and the corresponding pair in $\mathsf{Rel}(\mathcal{D}_s)(R)$ consists of the left and right marginals of μ.

Now, we can instantiate the premise of the rule for cobisimulations for this relation lifting on a coalgebra $\gamma \colon X \to \mathcal{D}_s(X)$ to obtain the following rule:

$$\frac{\forall \mu \in \mathcal{D}_s(\overline{R}).\, \gamma(x) \neq \mathcal{D}_s\pi_1(\mu) \vee \gamma(y) \neq \mathcal{D}_s\pi_2(\mu)}{x \mathrel{R} y} \tag{5.13}$$

However, this quantifies over the infinite set of subdistributions on the complement of R. This is rather unfortunate since we would like to prove apartness by giving a finite witness, as was the case for LTSs.

If we consider the example in (5.14), we see that x and y are cobisimilar, as they have different transition weights to sets of equivalent states. To see this using the rule (5.13), we could try to reason that given a distribution $\mu \in \mathcal{D}_s(\overline{R})$, the weight of the pair (x_1, y_1) would need to simultaneously match the transition probabilities from x to x_1 and from y to y_1, as x_1 and y_1 are only bisimilar to each other in this system. This is clearly not possible, so indeed all $\mu \in \mathcal{D}_s(\overline{R})$ satisfy the condition of rule (5.13).

(5.14)

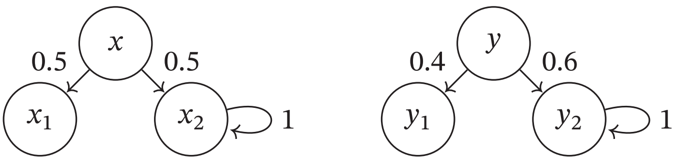

◂

We see two main issues with rule (5.13): the reasoning provided (where we choose the pair (x_1, y_1)) is not well reflected in the proof system; and the rule requires reasoning about an infinite set even for a simple *finite* coalgebra. This motivates much of the coming work, in which we move to an apartness notion defined dually to behavioural equivalence rather than coalgebraic bisimulation. The corresponding proof system will both exhibit the desired existential reasoning, and allow an optimisation giving finite proofs of apartness for both finite and infinite systems.

In the next section, we show the dualisation of behavioural equivalence, before instantiating to examples which will already illustrate the benefit of this approach over cobisimilarity. In Section 5.4, we go on to present the optimised proof rule, and show how this is beneficial by application to the examples of $\mathcal{D}_S$-coalgebras and stream systems.

5.3 Behavioural Apartness

In this section, we use a characterisation of behavioural equivalence based on so called *precongruences* due to Aczel & Mendler [AM89] to present a basic proof system for the dual: *behavioural apartness.* We will further apply the proof system to coalgebras for the subdistribution functor to show its benefits over the system provided in Section 5.2. Throughout this paper we work with coalgebras for an endofunctor $B\colon \mathsf{Set} \to \mathsf{Set}$.

5.3.1 Behavioural Equivalence

We start with the general definition of behavioural equivalence on Set coalgebras, based on cospans of coalgebra homomorphisms.

Definition 5.3.1: States x, y of a coalgebra $\gamma\colon X \to B(X)$ are said to be behaviourally equivalent if there is a coalgebra $\delta\colon Y \to B(Y)$ and coalgebra morphisms $f, g\colon (X, \gamma) \to (Y, \delta)$, such that $f(x) = g(y)$. ◂

Now the following definition (from [AM89]) will give an alternative characterisation of behavioural equivalence which is amenable to dualisation:

Definition 5.3.2: A precongruence on a coalgebra $\gamma\colon X \to B(X)$ is a relation $R \subseteq X \times X$ satisfying the following rule:

$$\frac{x\ R\ y}{Bq_R(\gamma(x)) = Bq_R(\gamma(y))} \tag{5.15}$$ ◂

Following Aczel & Mendler we call a precongruence which is an equivalence relation a *congruence*. They further showed that R is a congruence if and only if it is the kernel of a coalgebra morphism, so that the maximal (pre)congruence is equal to behavioural equivalence (also see [Gum99]). The following lemma tells us that the operator (implicit in Definition 5.3.2) mapping a relation R to the kernel $\ker(Bq_R \circ \gamma)$ is monotone. Therefore, behavioural equivalence can be obtained as its greatest fixed point.

Lemma 5.3.3 ([AM89, Lemma 4.1(2)]): For equivalence relations $R, S \subseteq X \times X$, we have $R \subseteq S \implies \forall x, y \in X.\, [Bq_R(\gamma(x)) = Bq_R(\gamma(y)) \implies Bq_S(\gamma(x)) = Bq_S(\gamma(y))]$. ◂

For further results relating such notions of equivalence see, e.g., [BSV04; Sta11] (the latter going beyond the category Set).

5.3.2 Dualising Behavioural Equivalence

The rule (5.15) can now be dualised to obtain a definition of behavioural apartness.

Definition 5.3.4: *Behavioural apartness* is the smallest relation satisfying the following rule:

$$\frac{Bq_{\overline{R}}(\gamma(x)) \neq Bq_{\overline{R}}(\gamma(y))}{x\ R\ y} \tag{5.16}$$

where $\overline{R}$ is the complement of the relation R. ◂

Dually to the case of precongruences and behavioural equivalence, this smallest relation will be an *apartness relation*.

Proposition 5.3.5: Behavioural apartness is an apartness relation, i.e., it is irreflexive ($\forall x \in X.\, \neg(x\ R\ x)$), symmetric, and cotransitive ($\forall x_1, x_2, x_3 \in X.\, (x_1\ R\ x_2 \implies x_1\ R\ x_3 \vee x_2\ R\ x_3)$). ◂

From definition to rule

In turning the rule of Definition 5.3.4 into a proof system, we make a number of changes. The use of the quotient map in the definition of the rule would require one writing a proof to ensure that the collection of pairs proved apart and used in a proof step forms an equivalence relation. To avoid this requirement, we allow an arbitrary relation R, with the only requirement that it consists of pairs which have been proved apart. Then, we introduce the following definition to replace the inequation $Bq_{\overline{R}}(\gamma(x)) \neq Bq_{\overline{R}}(\gamma(y))$:

$$t_1 \not\simeq^B_R t_2 := Bq_{\overline{R^s}}(t_1) \neq Bq_{\overline{R^s}}(t_2) \tag{5.17}$$

where R^s is the symmetric closure of R. Note that we are also taking the equivalence closure of $\overline{R^s}$, as in Definition 5.3.2 and also that symmetric closure does not commute with complement, i.e., taking the symmetric closure first really gives a distinct inequality. Our motivation for taking this symmetric closure will become clear when we come to examples, where it saves us the work of proving symmetric pairs apart.

The following lemma shows that we can also view this inequality directly in terms of apartness relations.

Lemma 5.3.6: For a relation $R \subseteq X \times X$, let R° denote the largest apartness relation contained in R (also called the apartness interior). Then

$$e(\overline{R}) = \overline{R^{\circ}} \tag{5.18}$$ ◂

We could thus equivalently define

$$t_1 \not\approx^R_B t_2 := Bq_{\overline{(R^s)^{\circ}}}(t_1) \neq Bq_{\overline{(R^s)^{\circ}}}(t_2) \tag{5.19}$$

wherein the implicit application of equivalence closure to $\overline{(R^s)^{\circ}}$ does not change the relation, as it is an idempotent operation.

Taking either definition, we propose the following proof rule

$$\frac{\forall (x', y') \in R.\, x' \mathrel{\#} y' \qquad \gamma(x) \not\approx^B_R \gamma(y)}{x \mathrel{\#} y} \tag{5.20}$$

Soundness and Completeness

We finish this section with a proof of soundness and completeness of the above rule, i.e., that states are behaviourally apart if and only if we can prove this using the rule. The theorem requires the additional assumption that the functor B is finitary. Some intuition for this condition can be gained from considering the final sequence of B and the cone of maps into it from the state space of a coalgebra $\gamma : X \to BX$ induced by the coalgebra structure.

This cone, which we will write as $(f_i : X \to B^i 1)_{i \in I}$, is defined by (transfinite) induction, starting from the unique map $f_0 : X \to 1$ given by finality. Then $f_{i+1} := Bf_i \circ \gamma$ for successor ordinals, and $f_\alpha := \lim_{i<\alpha} f_i$ for limit ordinals. By considering the kernels of these maps, we obtain a notion of n-step behavioural equivalence, i.e., states x, y are n-step behaviourally equivalent if $f_n(x) = f_n(y)$. Dually, states are n-step behaviourally apart if they are distinguished by f_n.

The crux of our soundness and completeness proof is then in showing that finite proof trees built from our proposed rule correspond to this notion of finite step apartness. The assumption that the functor B is finitary will then ensure that all states which are apart can be distinguished at some finite level, by results of Worrell [Wor05].

Theorem 5.3.7: Let B be a finitary endofunctor on Set, and $\gamma : X \to B(X)$ a coalgebra. Then states $x, y \in X$ are behaviourally apart if and only if we have a proof tree of finite height built from the following rule:

$$\frac{\forall (x', y') \in R.\, x' \mathrel{\#} y' \qquad \gamma(x) \not\approx^B_R \gamma(y)}{x \mathrel{\#} y} \tag{5.21}$$

◂

Note that the given "rule" is, strictly speaking, a family of rules indexed by relations $R \subseteq X \times X$. Thus, a proof tree may contain different instances of the rule, involving different choices of R.

Proof. The if direction holds by induction, as then the relation R is contained in behavioural apartness so that $\overline{R^s}$ contains behavioural equivalence. Using Lemma 5.3.3, we then see that $\gamma(x) \not\sim^B_R \gamma(y) \implies \gamma(x) \not\sim^B_{\underline{\#}} \gamma(y)$, where we write $\underline{\#}$ for behavioural apartness here. By Definition 5.3.4, we thus have $x \mathrel{\underline{\#}} y$.

For the other direction, we use the cone of functions $(f_i : X \to B^i 1)_{i \in I}$ defined earlier. We will first show that for any $n < \omega$ and pair of states $(x, y) \in \overline{\ker(f_n)}$, there is a proof tree of height n with $x \mathrel{\#} y$ as its root. The base case is trivial, as $\overline{\ker(f_0)}$ is empty. Now suppose that the property holds for some $N \geq 0$, and further that we have some $(x, y) \in \overline{\ker(f_{N+1})}$. This means (by definition) that

$$Bf_N(\gamma(x)) \neq Bf_N(\gamma(y)) \tag{5.22}$$

We then need to show that there is some set of pairs $R \subseteq X \times X$ so that

$$Bq_{\overline{R^s}}(\gamma(x)) \neq Bq_{\overline{R^s}}(\gamma(y)) \tag{5.23}$$

and we claim that $\overline{\ker(f_N)}$ is such an R, and that it gives us a proof tree of height N with $x \mathrel{\#} y$ as root.

By the induction hypothesis, for any $(x', y') \in \overline{\ker(f_N)}$, we have proof trees with $x' \mathrel{\#} y'$ at the root, and height N. Further, the premise (Eq. (5.23)) holds by definition of the kernel. Combining these gives us the required proof tree.

For the limit at ω, we have

$$\overline{\ker(f^\omega)} = \overline{\left(\bigcap_{i<\omega} \ker(f^i)\right)} = \bigcup_{i<\omega} \overline{\ker(f^i)} \tag{5.24}$$

so that for $(x, y) \in \overline{\ker(f^\omega)}$, there is some $i < \omega$ with $(x, y) \in \overline{\ker(f^i)}$. We obtain the required proof tree by (transfinite) induction.

It is further shown by Worrell [Wor05], that the final sequence for finitary functors on Set stabilises at $\omega 2$ and, more importantly here, that the connecting maps $B_\beta \to B_\lambda$ for $\omega \leq \lambda \leq \beta \leq \omega 2$ are injective. Thus, for $\omega \leq \beta \leq \omega 2$, we have $f^\beta(x) \neq f^\beta(y) \implies f^\omega(x) \neq f^\omega(y)$, again giving the required proof tree. ∎

We note that more general versions of the above could perhaps be developed using results from [HKC18]. Namely, results showing that coinductive predicates for finitary functors can be constructed via a final sequence which stabilises after ω steps. Dualising, there should be a correspondence between the inductively defined set of proof trees, and the inductive predicate behavioural apartness. In the current work, we focus on behavioural equivalence on coalgebras in Set, and are thus able to use the earlier results due to Worrell.

Before we show how rule (5.21) can be improved further, we apply it to the coalgebra for the subdistribution functor of Example 5.2.7 to show how it improves on rule (5.2) of Section 5.2.

5.3.3 Example: Subdistributions

We saw in Section 5.2 how the definition of cobisimilarity for coalgebras of the subdistribution functor does not allow us to easily produce proofs. Here, we show that our new notion of behavioural apartness is more usable in this regard.

For this, we first elaborate how the functor $\mathcal{D}_s$ acts on the quotient map of an equivalence relation, i.e., a map $q_R \colon X \twoheadrightarrow X/e(R)$ mapping an element to its equivalence class under the equivalence closure $e(R)$ of R. The image of such a map under $\mathcal{D}_s$ is given by

$$\mathcal{D}_s(q_R)(\mu)([z]_R) = \sum_{x \in [z]_R} \mu(x) \tag{5.25}$$

where, by $[z]_R$ we mean the equivalence class of z under the equivalence closure $e(R)$ of R. We may also write this sum as $\mu[z]_R$, an instance of the general notation $\mu(S \subseteq X) = \sum_{s \in S} \mu(s)$.

Then rule (5.21), for some coalgebra $\gamma \colon X \to \mathcal{D}_s(X)$, instantiates to

$$\frac{\forall (x', y') \in R.\, x' \# y' \qquad \exists z \in X.\, \gamma(x)[z]_{\overline{R^s}} \neq \gamma(y)[z]_{\overline{R^s}}}{x \# y} \tag{5.26}$$

Returning to the earlier concrete example of (5.14), we can now produce a proof which closely matches the reasoning we suggested earlier. We start by noting that for states such as x_1 and y_2 with one state having no outgoing probability weight, we have the following proof of behavioural apartness:

$$\frac{R = \varnothing \qquad \gamma(x_1)[y_2]_{\top_X} = 0 \neq 1 = \gamma(y_2)[y_2]_{\top_X}}{x_1 \# y_2} \tag{5.27}$$

where $\top_X$ is the total relation on X. In this same way, we can prove $x' \# y'$ for all pairs in the set $S_1 = \{(x_1, x_2), (x_1, y_2), (x_2, y_1), (y_1, y_2)\}$. We go on to prove that both x and y are apart from both of x_1, y_1. For example, we can prove

$$\frac{R = S_1 \qquad \gamma(x)[x]_{\top_X} = 1 \neq 0 = \gamma(y_1)[x]_{\top_X}}{x \mathrel{\#} y_1} \tag{5.28}$$

Although we have already proved a number of apartness pairs, taking the apartness interior still yields the empty set. This is forced by the cotransitivity axiom: for example as $(x_1, x_2) \in S_1$, we must prove all other states apart from at least one of x_1 and x_2 if the pair (x_1, x_2) is to be in the apartness interior, which we have not yet done. This means we can so far only distinguish states for which the total outgoing weight is different.

Only once we have proven apartness for all pairs in the set

$$S := S_1 \cup \{(x, x_1), (x, y_1), (y, x_1), (y, y_1)\} \tag{5.29}$$

can we give the proof of $x \mathrel{\#} y$:

$$\frac{R = S \qquad \gamma(x)[x_1]_{\overline{R^s}} = 0.5 \neq 0.4 = \gamma(y)[x_1]_{\overline{R^s}}}{x \mathrel{\#} y} \tag{5.30}$$

Note that here

$$\overline{R^s} = e(\overline{R^s}) \tag{5.31}$$

$$= \{(x, y), (x_1, y_1), (x_2, y_2), (x, x_2), (y, y_2), (x, y_2), (y, x_2)\}^s \cup \Delta_X \tag{5.32}$$

This demonstrates that we are indeed able to provide a proof of behavioural apartness using rule (5.21) which is finite and built up from inequations exhibiting differences in transition weights. The situation can however still be improved, by reducing the work required in our proofs. As mentioned, we were required to prove for instance the apartness pair $x \mathrel{\#} y_1$ to make the apartness interior non-trivial. As x does not contribute to the transition distributions of the states we are proving apart, this should not be necessary for the proof. Such proof obligations arise because we take the apartness interior/equivalence closure with respect to the entire state space, and (co)transitivity forces that we include many “unnecessary” pairs. In the next section, we show that we can restrict to apartness relations on just those states that are in the supports of the successor distributions, or more generally, states which are reachable in one step.

5.4 Optimised Rules for Behavioural Apartness

To further improve our proof system, we show that it is sound and complete to prove apartness using an inductive step in which apartness need only be proved on states which are “reachable in one step” from the states which we are proving apart. We make this notion precise in the following definition.

Definition 5.4.1: Given a coalgebra $\gamma\colon X \to B(X)$, we will say that a subobject $m\colon Z \rightarrowtail X$ is a one-step covering of a subobject $s\colon S \rightarrowtail X$ if there exists a map $g\colon S \to BZ$ such that $Bm \circ g = \gamma \circ s$. We may also say that Z (one-step) covers S via g.

Further, we denote the set of one-step coverings of a subobject S by $\varDelta_S$. ◂

Using this definition, we show that inequality of pairs of states under the map $Bq_{\overline{R}} \circ \gamma$ which we have used in all our rules so far, is equivalent to inequality under the map $Bq_{\overline{R}|_Z} \circ g$ where we restrict the relation R to a Z that is a one-step covering of some S via g.

Lemma 5.4.2: Let $\gamma\colon X \to B(X)$ be a Set-coalgebra, $R \subseteq X \times X$ an equivalence relation, and $S \rightarrowtail X$ a subobject. For any non-empty $m\colon Z \rightarrowtail X$ which is a one-step covering of S via $g\colon S \to B(Z)$, we have

$$Bq_R(\gamma(x)) \neq Bq_R(\gamma(y)) \iff Bq_{R|_Z}(g(x)) \neq Bq_{R|_Z}(g(y)) \tag{5.33}$$

where $q_R\colon X \to X/R$, $q_{R|_Z}\colon Z \to Z/R|_Z$ are quotient maps on X and Z respectively. ◂

This means, in other words, that we can determine behavioural apartness by looking only at apartness on one-step coverings. The given definition is a generalised version of the notion of the *base of a functor* originally due to Blok [Blo12] which can be stated as the smallest one-step covering. It has been shown that this notion indeed instantiates to one-step reachability [BKR19; Wiß+20]. As an example, for a coalgebra $\gamma\colon X \to \mathcal{D}_s(X)$ for the subdistribution functor, the base for a singleton $\{x\} \rightarrowtail X$ is exactly the support of the successor distribution of x. More interestingly for our applications are one-step coverings of pairs of states x, y, which will instantiate to the union of the supports of x and y in the subdistribution case. In general, the base may not always exist. Indeed, it requires the base category to be complete and well-powered, and the behaviour functor to preserve wide intersections (see [BKR19, Prop. 12]). The restriction to finitary functors on Set (which we make in Theorem 5.3.7) however implies these conditions. Despite this, we choose not to take the base itself in the rule which we give in Theorem 5.4.3; instead we allow arbitrary one-step coverings, to keep our results more general.

Proof of Lemma 5.4.2. Suppose we have a map g such that $\gamma \circ s = Bm \circ g$, with $m\colon Z \rightarrowtail X$ non-empty. We form the restriction of a relation to the subobject $m\colon Z \rightarrowtail X$ by the following pullback:

$$
\begin{tikzcd}
R|_Z \arrow[r, tail] \arrow[d, tail] \arrow[dr, phantom, "\lrcorner", very near start] & Z \times Z \arrow[d, tail, "m\times m"] \\
R \arrow[r, tail] & X \times X
\end{tikzcd}
\tag{5.34}
$$

We then obtain the following diagram:

$$
\begin{tikzcd}
R|_Z \arrow[r, shift left, "\pi_1^Z"] \arrow[r, shift right, "\pi_2^Z"'] \arrow[d] & Z \arrow[r, two heads, "q_R|_Z"] \arrow[d, tail, "m"] & Z/R|_Z \arrow[d, dashed, "!"] \\
R \arrow[r, shift left, "\pi_1^X"] \arrow[r, shift right, "\pi_2^X"'] & X \arrow[r, two heads, "q_R"] & X/R
\end{tikzcd}
\tag{5.35}
$$

where $!\colon Z/R|_Z \to X/R$ is the unique map which arises from the fact that $q_R \circ m$ coequalizes π_1^Z, π_2^Z by commutativity of the left square. We now consider the commutative diagram:

$$
\begin{tikzcd}
S \arrow[r, tail, "s"] \arrow[dr, "g"'] & X \arrow[r, "\gamma"] & B(X) \arrow[rr, "Bq_R"] & & B(X/R) \\
& BZ \arrow[ur, "Bm"'] \arrow[rr, "Bq_{(R)|_Z}"'] & & B(Z/(R)|_Z) \arrow[ur, "B!"'] &
\end{tikzcd}
\tag{5.36}
$$

This means that $Bq_R \circ \gamma \circ s = Bq_R \circ Bm \circ g = B! \circ Bq_{(R)|_Z} \circ g$ so that

$$
\begin{aligned}
& Bq_R(\gamma(s(x))) \neq Bq_R(\gamma(s(y))) && (5.37)\\
\iff & B! \circ Bq_{(R)|_Z} \circ g(x) \neq B! \circ Bq_{(R)|_Z} \circ g(y) && (5.38)\\
\implies & Bq_{(R)|_Z} \circ g(x) \neq Bq_{(R)|_Z} \circ g(y) && (5.39)
\end{aligned}
$$

Now, if $B!$ is mono, the last implication becomes a bi-implication, which is what we want. In Set this holds as follows: we have $!([z]_{R|_Z}) = [z]_R$ and it is clear that $[z]_R = [z']_R \implies [z]_{R|_Z} = [z']_{R|_Z}$, i.e., ! is mono. In fact, it is *split* mono, as it is a function with non-empty domain by the assumption that Z is non-empty. Split monos are preserved by all functors, so that $B!$ is mono as required. ∎

We are now in a position to give our final proof rule for behavioural apartness, so that we can prove behavioural apartness $x \# y$ with an inductive step, where only pairs of states in a one-step covering of $\{x, y\}$ need to be proved apart.

A possible choice in the definition of the family of rules, is to have a rule for every one-step covering of $\{x, y\}$. This would include the rule involving the whole state space X, reducing to the earlier defined set of rules. While this choice makes the proof of completeness trivial, we can do better. It is possible to first fix a choice of one-step covering for each $\{x, y\}$, and then restrict to the rules which involve these coverings only. In particular, in case the base exists, this suffices.

Theorem 5.4.3: Let $\gamma : X \to B(X)$, and let $\Delta \subseteq \bigcup_{x,y \in X} \Delta_{\{x,y\}}$ be a subset of the one-step covers of sets $\{x, y\}$ such that for every pair $x, y \in X, \Delta_{\{x,y\}} \cap \Delta \neq \emptyset$. Then the family of rules defined as follows, with Z ranging over Δ and $R \subseteq Z \times Z$, is sound and complete for behavioural apartness on γ:

$$\frac{\forall (x', y') \in R.\, x' \# y' \qquad g(x) \not\sim^B_{R,Z} g(y)}{x \# y} \tag{5.40}$$

where we define

$$t_1 \not\sim^B_{R,Z} t_2 \;:=\; Bq_{e(\overline{R^s})|_Z}(t_1) \neq Bq_{e(\overline{R^s})|_Z}(t_2) \tag{5.41}$$ ◀

Note that we omit the subscript on the one-step covering when it is clear from context and now explicitly take the equivalence closure before restricting to Z. The latter is required for the application of Lemma 5.4.2 in the following proof.

Proof. We prove soundness by induction on the structure of a proof tree with root $x \# y$ built using the above rule involving a relation R. By Lemma 5.4.2 instantiated to the relation $e(\overline{R^s})$, we see that $\gamma(x) \not\sim^B_R \gamma(y)$ will hold, so that the condition of the sound rule (5.21) holds.

For completeness, the proof is essentially the same as for Theorem 5.3.7, but we take R to be $\ker(f_N)|_Z$ in the induction step. Then we have $\ker(f_N)|_Z = \ker(f_N \circ m)$. After some rewriting, the inequality $g(x) \not\sim^B_{R,Z} g(y)$ becomes $Bq_{\ker(f_N \circ m)}(g(x)) \neq Bq_{\ker(f_N \circ m)}(g(x))$, which holds by definition of the quotient map of the kernel. ∎

Example 5.4.4 (Subdistributions): We return to the example of coalgebras for the subdistribution functor, and show how the new rule (5.40) improves on rule (5.21).

First, note that we can specialise the rule to subdistributions:

$$\frac{\forall (x', y') \in R.\, x' \# y' \qquad \exists z \in Z.\, g(x)[z]_E \neq g(y)[z]_E}{x \# y} \tag{5.42}$$

where we choose $Z := \mathrm{supp}(\gamma(x)) \cup \mathrm{supp}(\gamma(y))$ (so that $g : \{x, y\} \to Z$ is the restriction of γ to just the states x and y) and let $E := e(\overline{R^s})|_Z$. We will again prove the apartness $x \mathrel{\#} y$ in the example of (5.14), using our new rule. For the last proof step, we have

$$\frac{\begin{array}{c} R = \{(x_1, y_2), (x_1, x_2), (y_1, x_2), (y_1, y_2)\} \\ g(x)[x_1]_E = 0.5 \neq 0.4 = g(y)[x_1]_E \end{array}}{x \mathrel{\#} y} \tag{5.43}$$

where $E = \{(x_1, y_1), (x_2, y_2)\} \cup \Delta_Z$.

For the proof obligations given by the relation R we give one pair as an example; the rest are similar.

$$\frac{R = \varnothing \qquad g(x_1)[y_2]_E = 0 \neq 1 = g(y_2)[y_2]_E}{x_1 \mathrel{\#} y_2} \tag{5.44}$$

We see that while the proof step for such leaves does not change much when using the rule of Theorem 5.4.3, the proof in its totality is easier to provide than what we showed in Section 5.3.3. It also fits better with the desired reasoning based on supplying witnesses of apartness and only reasoning about successors. ◀

Example 5.4.5 (Streams): Taking $B = \Sigma \times (-)$ for some set of symbols Σ, we call B-coalgebras *stream systems*, and instantiate the rules (5.21) and (5.40) to this system type. We write a stream system as $\langle o, t\rangle \colon X \to \Sigma \times X$, so that the condition in (5.21) can be rewritten as follows:

$$\mathrm{id}_\Sigma \times q_{\overline{R^s}}(\langle o, t\rangle(x)) \neq \mathrm{id}_\Sigma \times q_{\overline{R^s}}(\langle o, t\rangle(y)) \tag{5.45}$$
$$\iff (o(x), q_{\overline{R^s}}(t(x))) \neq (o(y), q_{\overline{R^s}}(t(y))) \tag{5.46}$$
$$\iff o(x) \neq o(y) \lor \neg(t(x)\ e(\overline{R^s})\ t(y)) \tag{5.47}$$
$$\iff o(x) \neq o(y) \lor t(x)\ (R^s)^\circ\ t(y) \tag{5.48}$$

In words, two states of a stream system behave differently if they have different outputs (the heads of the streams they represent are different) or their successors behave differently (the tails of the streams they represent are different).

$$\frac{o(x) \neq o(y)}{x \mathrel{\#} y} \qquad \frac{\forall (x', y') \in R.\, x' \mathrel{\#} y' \qquad t(x)\ (R^s)^\circ\ t(y)}{x \mathrel{\#} y} \tag{5.49}$$

In this form, the right-hand rule involving successors requires us to take the interior. As we have seen in Section 5.3.3, this will be empty unless we have proven enough apartness pairs. Consider, for instance, the following simple example

$$x_0 \xrightarrow{1} x_1 \xrightarrow{1} x_2 \xrightarrow{2} x_3 \xrightarrow{3} x_4 \xrightarrow{5} \cdots \tag{5.50}$$

in which $o(x_i)$ is the i-th Fibonacci number, so that x_0 generates the full stream of Fibonacci numbers and x_1 generates the stream of Fibonacci numbers excluding the first one. It should be clear that these states are therefore behaviourally apart. However, in order to prove this with the above rules we must use the rule involving successors, and hence provide a useful relation R. For this R to be an apartness on the whole state space and to be non-empty, we must already prove an infinity of apartness pairs (at least one including each of the states due to cotransitivity).

The solution is of course provided by our optimised rule (5.40). For a stream system $\langle o, t \rangle \colon X \to \Sigma \times X$, we can choose a rather simple one-step covering for each pair x, y, namely $\{t(x), t(y)\}$. Then, the only choice of R yielding a usable rule (up to symmetry), is $\{(t(x), t(y))\}$, so that rule (5.21) instantiated to stream systems can be simplified to the following, in which we split the disjunction in the premise into two rules:

$$\frac{o(x) \neq o(y)}{x \mathrel{\#} y} \qquad \frac{t(x) \mathrel{\#} t(y)}{x \mathrel{\#} y} \tag{5.51}$$

For the example (5.50), the one-step covering of $\{x_0, x_1\}$ to be used is thus $Z = \{x_1, x_2\}$ with the map $g \colon \{x_0, x_1\} \to \Sigma \times \{x_1, x_2\}$ defined by $g(x_0) = (1, x_1)$ and $g(x_1) = (1, x_2)$. Then the following

$$\frac{\dfrac{o(x_1) \neq o(x_2)}{t(x_0) = x_1 \mathrel{\#} x_2 = t(x_1)}}{x_0 \mathrel{\#} x_1} \tag{5.52}$$

is a finite proof of apartness of x_0 and x_1. ◂

5.4.1 Inductive Characterisation of $\not\asymp$

In [GJ21, Appendix A], an inductive definition of cobisimilarity for Kripke polynomial functors is given. Here, we show the instantiation of the relation $\not\asymp_R^B$ to a class of functors extending the Kripke polynomial functors. This will allow us to easily obtain proof rules for coalgebras of functors in this class.

We consider a slight restriction of Kripke polynomial functors on Set as defined by Jacobs [Jac16] extended with the subdistribution functor, the syntax of which we give using the following grammar:

$$B ::= \mathsf{Id} \mid A \mid B \times B \mid B + B \mid B^A \mid \mathcal{P}_\alpha B, \alpha \in \mathcal{N} \mid \mathcal{D}_s B \tag{5.53}$$

where A is any set and $\mathcal{N}$ is the class of cardinals. The restriction is to binary coproducts, which matches the presentation in [GJ21, Appendix A]. The characterisation of $\not\asymp_R^{\mathcal{P}_\alpha B}$ does not actually depend on the cardinal α, but we will only have a complete proof system if we choose $\alpha = \aleph_0$ (i.e., the finite powerset functor) as shown in Theorem 5.3.7.

We can now instantiate the definition of the relation $\not\asymp^B_R$ given in (5.17), with the functor B built from the above grammar. We do this for $\not\asymp^B_R$ here to ease notation. To replace $\not\asymp^B_R$ with $\not\asymp^B_{R,Z}$ we must take the apartness interior in the first statement with respect to Z, and the complement in the case $B = \mathcal{D}_s B_1$ must also be taken with respect to Z. More precisely, we have that $\overline{\not\asymp^{B_1}_{R,Z}} = (Z \times Z) \setminus \not\asymp^{B_1}_{R,Z}$. The proof in each case is a routine calculation.

Lemma 5.4.6: For a relation $R \subseteq X \times X$ and functors $B, B_1, B_2 \colon \mathsf{Set} \to \mathsf{Set}$ we have the following inductive characterisation of $t_1 \not\asymp^B_R t_2$.

- If $B = \mathsf{Id}$, then $t_1 \not\asymp^B_R t_2 \iff t_1\ (R^s)^\circ\ t_2$.
- If $B = A$, then $t_1 \not\asymp^B_R t_2 \iff t_1 \neq t_2$.
- If $B = B_1 \times B_2$, then $(u_1, v_1) \not\asymp^B_R (u_2, v_2) \iff u_1 \not\asymp^{B_1}_R u_2 \vee v_1 \not\asymp^{B_2}_R v_2$.
- If $B = B_1 + B_2$, then $t_1 \not\asymp^B_R t_2 \iff [t_1, t_2 \in B_1 X \implies t_1 \not\asymp^{B_1}_R t_2] \wedge [t_1, t_2 \in B_2 X \implies t_1 \not\asymp^{B_2}_R t_2]$
- If $B = B_1^A$, then $t_1 \not\asymp^B_R t_2 \iff \exists a \in A.\, t_1(a) \not\asymp^{B_1}_R t_2(a)$
- If $B = \mathcal{P}_\alpha B_1$, then $t_1 \not\asymp^B_R t_2 \iff [\exists u \in t_1.\, \forall v \in t_2.\, u \not\asymp^{B_1}_R v] \vee [\exists v \in t_2.\, \forall u \in t_1.\, u \not\asymp^{B_1}_R v]$
- If $B = \mathcal{D}_s B_1$, then $t_1 \not\asymp^B_R t_2 \iff \exists z \in \mathcal{B}_1 X.\, t_1[z]_{\overline{\not\asymp^{B_1}_R}} \neq t_1[z]_{\overline{\not\asymp^{B_1}_R}}$ ◂

As discussed by Sokolova in [Sok11], the (sub)distribution functor can be seen as an instance of the finitely supported monoid valuations functor in a set S, denoted by $\mathcal{M}^{-}_{S,f}$. The valuations given by this functor, are maps $v \colon X \to M$ for a monoid M with finite support, such that $\sum_{x \in X} v(x) \in S$. Taking M to be the non-negative reals with addition, and $S = \{1\}$, recovers the full distribution functor. Taking instead $S = [0, 1]$ recovers the subdistribution functor. This is also a generalisation of the finitely supported multiset functor by taking M the natural numbers with addition, and $S = \mathbb{N}$ (see also [Gum09; GS01]). Due to the construction of our proof system, rules for such a functor can be straightforwardly obtained by calculation of $\not\asymp^{\mathcal{M}^{-}_{S,f}}_R$.

Example: MDPs

There are of course now many examples for which we obtain a proof system. Here, we take coalgebras for the functor $\mathcal{D}_s(-)^A$, with A some finite set of actions, which we choose to call Markov Decision Processes (without rewards).[1] We consider the following MDP:

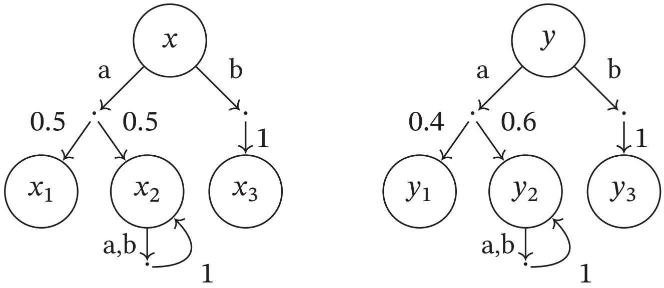

States with no outgoing edge are those s for which $\gamma(s)(a) = 0$ (the zero distribution) for all $a \in A$. To instantiate the generic rule (5.40) to this setting, we elaborate the involved inequality:

$$\gamma(x) \not\leq_R^{\mathcal{D}_s^A} \gamma(y) \tag{5.54}$$

$$\iff \exists a \in A.\, \gamma(x)(a) \not\leq_R^{\mathcal{D}_s} \gamma(y)(a) \tag{5.55}$$

$$\iff \exists a \in A.\, \exists z \in Z.\, \gamma(x)(a)[z]_{e(\overline{R^s})|_Z} \neq \gamma(y)(a)[z]_{e(\overline{R^s})|_Z} \tag{5.56}$$

where Z is any one-step covering for $\{x, y\}$. The rule is thus:

$$\frac{\forall (x', y') \in R.\, x' \# y' \qquad \exists a \in A.\, \exists z.\, \gamma(x)(a)[z]_{e(\overline{R^s})|_Z} \neq \gamma(y)(a)[z]_{e(\overline{R^s})|_Z}}{x \# y} \tag{5.57}$$

We now take $\sigma = \mathrm{a}$, the one-step covering $Z = \{x_1, x_2, x_3, y_1, y_2, y_3\}$, and

$$R = \{(x_1, x_2), (x_2, x_3), (y_1, y_2), (y_2, y_3), (x_1, y_2), (x_2, y_3), (x_3, y_2), (x_2, y_3)\} \tag{5.58}$$

Then, for example, $z = x_2$ gives us

$$\frac{\forall (x', y') \in R.\, x' \# y' \qquad \gamma(x)(a)[x_2]_{e(\overline{R^s})|_Z} = 0.5 \neq 0.6 = \gamma(y)(a)[x_2]_{e(\overline{R^s})|_Z}}{x \# y} \tag{5.59}$$

---

[1]The name MDP is perhaps a little overloaded. Such systems have also been called Labelled Markov Processes, but we choose not to use that name here to avoid confusion with the coming chapter

5.5 Conclusion

There are a number of avenues for future work. The first is based on the established link between notions of bisimilarity/behavioural equivalence (and their duals) and (coalgebraic) modal logic [Sch08; FKP17; KMS20; WMS22; Geu22]. We would like to investigate how proofs of behavioural apartness in our system relate to formulas in a corresponding logic. Namely, can we extract formulas distinguishing states which are behaviourally apart? Such results already seem close by if we are able to choose our logic to match the reasoning present in our proof system. Indeed, a proof of behavioural apartness seems close to a proof that a formula is true in one state and its negation is true in the other.

A more recent notion of equivalence based on *codensity lifting* [KSU18; Spr+21; Kom+22] could also be an approach to more general notions of inequivalence. It has been shown how this can be used to define, e.g., behavioural preorders such as simulation [KSU18] and also quantitative equivalences, which have further been linked to quantitative logics [Kom+21; KM18; Gon+23; For+23; Beo+23; Beo+24]. We would like to investigate whether the dual notions (which we may call codensity apartness) allow an easier development of expressive logics in these settings. Part of this line of research should be to apply the ideas of Section 5.4 to general liftings beyond that for behavioural apartness (also beyond Set). For instance, can we prove codensity apartness while only looking at states reachable in one step? We hope this would simplify the development of proof systems in the quantitative setting.

In the next chapter, we will move a step closer to these goals by developing a derivation system for lower bounds on the behavioural distance between states in systems with probabilistic transitions and labels on states. We will also relate proofs in this system to formulas in a logic with a quantitative semantics.

CHAPTER 6

Constructing Witnesses for Lower Bounds on Behavioural Distances

Testaments to nous and short
Suzuka turns (9)

6.1 Introduction

Up to now, we have considered notions of equivalence, such as bisimilarity, which are *qualitative* in the sense that they consider states of a transition system to be either equivalent or not; there are no degrees of equivalence. When studying systems involving probabilistic transitions, such qualitative definitions are usually considered too strict; states may be *inequivalent* or *distinguishable* under bisimilarity despite their behaviour being difficult to distinguish by an observer (this problem was first noted by Giacalone, Jou, and Smolka [GJS90]).

To better capture the (in)equivalence of states, quantitative notions of *behavioural distances/metrics* may be used [Des+99; Bre17; BW05]. These assign to each pair of states a number (e.g., in the interval $[0, 1]$) representing how close (or how far apart) their behaviours are. Determining these distances has been studied algorithmically, with procedures developed for approximating the distance [BSW08], and the exact computation of distances [CBW12;

Bac+17; Bac+21; Tan18]. The definition of behavioural distances as least or greatest fixed points (depending on the chosen ordering), also gives them a universal property yielding a (co)inductive proof principle [HJ98; Bal+18; BKP23]. The corresponding proofs are of bounds on the distances showing states to be equivalent to some degree. This is analogous to the qualitative proof technique of exhibiting some bisimulation containing a pair of states, thereby showing them to be bisimilar.

Apartness Orthogonally, interest in (qualitative) apartness of states has been growing as an inductive counterpart to bisimilarity [GJ21]. Rather than defining when states behave the same, apartness defines when there is some observable difference between them. A reason for interest in apartness is its inductive and potentially constructive nature. Indeed, this was the original motivation, going back to the school of Brouwer [Hey66]. In the setting of state-based systems, this means giving some (finite) evidence or witness for a difference in behaviours. Think of, for example, a word which is accepted by one state of a finite automaton but not another, or a difference in probability of making a certain observation in probabilistic systems such as labelled Markov chains or Markov decision processes.

As in the case of bisimilarity, we would like these notions to be as robust as possible, making behavioural distances a clear area of interest. In the quantitative setting, the dual of the existing coinductive proof methods allows us to obtain lower bounds on distances between states, i.e., we can show states to be at least a certain distance apart. This type of bound has recently been explored in [Bal+23]. They define a measure of how much a candidate for the least fixed point can be decreased. If no such decrease is possible, we have a lower bound.

In this work, we take an alternative approach, based essentially on Kleene's chain construction of least fixed points [Kle52]. For the case of behavioural distances, this starts from an order-preserving functional, say $\Gamma : \mathsf{PMet}_X \to \mathsf{PMet}_X$, on the space of pseudometric spaces on a set X. To approach the least fixed point $\mu\Gamma$ from below, we can start from the constant zero distance $\bot$ and iteratively apply Γ giving the chain $\bot \leq \Gamma(\bot) \leq \Gamma^2(\bot) \leq \dots$. As is noted in [Bal+23], fully applying Γ iteratively in this way is not a desirable means of obtaining bounds. Instead, we will develop an inductive derivation system allowing the construction of lower bounds for chosen pairs of states.

However, simply translating the definition of the behavioural distance to proof rules does not work; to obtain a usable derivation system we must show that proof steps need only consider direct successors of the involved states. This means that finite derivations can be constructed also for systems with infinite state space. We further reduce the work required for proofs, and thus also the size of the derivations, by showing that recursive proofs are only required for a subset of these successors.

The judgments proved in our system are of the form $x \#_\varepsilon y$ for x, y states in a labelled Markov chain (LMC) and ε a rational in the interval $[0, 1]$. To ensure a correct relation to the behavioural distance, we show two properties. First, soundness: that if we can prove $x \#_\varepsilon y$, then ε is a lower bound on the behavioural distance between x and y. Second, a form of completeness which we call *approximate* completeness. Usual completeness with respect to the behavioural distance would state that any distance between states can be proved in the derivation system. However, in the spirit of apartness, we consider *finite* evidence, which can only witness finite approximations of distances in general. We thus show that lower bounds can be derived with arbitrarily small error with respect to the true distance.

Logics Evidence of differences in behaviour can also be given in the form of formulas in some modal logic. This is closely related to Hennessy-Milner type theorems, which show for a given logic that bisimilarity and logical equivalence coincide. Such theorems have been shown for probabilistic bisimilarity and a logic with a modality parameterised by rational probabilities [DEP98; DEP02; FKP17]. Later, a quantitative analogue was shown by relating a real-valued logic introduced in [Des+99] to behavioural distances [Bre+07]. The most interesting part in the qualitative case is the inclusion of logical equivalence in bisimilarity, also called expressiveness, which can dually be shown by giving, for each pair of non-bisimilar states, a formula which distinguishes them. Quantitatively, this means giving a formula for which the difference in interpretations matches the behavioural distance as closely as possible.

This correspondence of behavioural distances and modal formulas has been investigated for labelled Markov chains (LMCs) in [RB23]. In their terminology, a construction is given of formulas "explaining" the distance between states. Due to the chosen logic, and the possibility of infinite behaviours in LMCs induced by loops, this can not be done exactly. Instead, it is shown that for any finite approximation $\Gamma^i(\bot)(x, y)$ of the distance of states as in the above chain, a formula can be constructed such that $|[\![\varphi]\!](x) - [\![\varphi]\!](y)|$ (the difference in interpretations on states x and y) is equal to the approximation.

We finish by relating our new derivation system to the work of [RB23] by showing that for any derivation, a formula in the modal logic of *op. cit.* can be constructed which witnesses the same bound. These witnessing formulas are an improvement as they can be given for infinite state systems, and they will be smaller in many examples. Further, for any formula a proof tree witnessing the same lower bound can be given whose depth will be equal to the modal depth of the formula. For more fine-grained notions of size counting the total number of steps in a derivation and number of operators in a formula, the derivations will be larger in general, as they are dependent on the system and thus contain more information. This is exactly what facilitates the aforementioned improvements; we see the steps which lead us to conclude a difference in behaviours which are otherwise somewhat hidden in the semantics of the logic. Proofs also focus on pairs of states, so that at each step we see which states of a system are being used to exhibit a lower bound.

Contributions Summarising the contributions of the chapter:

- We define an inductive derivation system for lower bounds on behavioural distances in labelled Markov chains (Section 6.3)
- We show the soundness and approximate completeness of the system with respect to the behavioural distance (Sections 6.3.1 and 6.3.2)
- We show a constructive correspondence between proofs in our system and formulas of a modal logic (Sections 6.4 and 6.5)

We illustrate all of the above with examples, including a system with an infinite state space which was not covered by this modal logic (Example 6.5.6).

Related Work The line of work focusing on proofs of apartness for state-based systems was (re)started by Geuvers and Jacobs [GJ21], with further work on the relation to distinguishing formulas in [Geu22]. A proof system for an apartness notion dual to coalgebraic behavioural equivalence was given in the previous chapter. The current chapter builds directly on these works.

The coinductive proof principle in the context of behavioural distances has been explored coalgebraically in, e.g., [Rut98; Bon+17; Bal+18; BKP23]. In the greatest fixed point characterisation, the ordering is reversed compared to the definition we use in the rest of this work. Coinduction thus leads to upper bounds on distances under our definition. An approach focused on bounding greatest fixed points from above (but which dually bounds least fixed points from below as we will do) has more recently been given in [Bal+23], as discussed above. There is however no construction given of formulas demonstrating proved bounds.

The construction of distinguishing formulas has also been studied in the qualitative setting in, e.g., [MG23; MG24; KMS20; WMS22]. We discuss these works further in Section 6.6.

A more general account of bounding distances from above is the area of quantitative equational theories [MPP16], which has been applied to give a calculus for upper bounds on distances in Markov chains [Bac+18] and regular expressions [Roz24].

Notation We will write $\mathcal{D}_{\mathbb{Q}}(X)$ for the set of finitely-supported rational distributions on the set X. These are maps $\mu \colon X \to [0,1] \cap \mathbb{Q}$ such that $\mathrm{supp}(\mu) := \{x \in X \mid \mu(x) \neq 0\}$ is finite and $\sum_{x \in X} \mu(x) = 1$. We may also write such distributions as formal sums: $\sum_{x \in X} \mu(x) \, |x\rangle$, with the $|x\rangle$ ("ket") notation taken from [Jac] and used simply to separate states and their associated probabilities. It can also be thought of as indicating a sum of Dirac distributions. From now on, we will write $[0,1]_{\mathbb{Q}}$ for $[0,1] \cap \mathbb{Q}$.

We denote by PMet_X the set of pseudometric spaces on a set X, i.e., pairs (X, d) with $d \colon X \times X \to [0,1]$ a pseudometric. We order the unit interval with the usual ordering of the reals, and pseudometrics inherit this ordering pointwise, so that $d_1 \leq d_2$ iff $\forall x, y \in X.\, d_1(x,y) \leq d_2(x,y)$. The smallest element $\bot$ is thereby the constant zero distance. Further, for two pseudometric spaces $(X, d_X) \in \mathsf{PMet}_X$ and $(Y, d_Y) \in \mathsf{PMet}_Y$ a map $f \colon (X, d_X) \to (Y, d_Y)$ will be assumed to be *non-expansive* (also called *1-Lipschitz*, *short*, etc.), i.e., $\forall x, y \in X.\, d_Y(f(x), f(y)) \leq d_X(x,y)$. The Euclidean distance is denoted $d_e \colon [0,1] \times [0,1] \to [0,1]$.

In the interest of space, we write $\mu \cdot h$ for $\sum_{x \in X} \mu(x) \cdot h(x)$ where $\mu \colon X \to [0,1]_{\mathbb{Q}}$ is a distribution and $h \colon X \to \mathbb{R}$ is an arbitrary function. This is motivated by viewing the distributions and functions as vectors indexed by their common domain, so that the operation is the vector dot product. In the sequel, we will often restrict h to maps into $[0,1]_{\mathbb{Q}}$.

6.2 Behavioural Distances on LMCs

We start with the definition of the type of system which we study in the remainder of the chapter: labelled Markov chains.

Definition 6.2.1: A labelled Markov chain (LMC) consists of the following data:

- a set of states X;
- a (non-empty) set of labels L;

- a (finitely branching) probabilistic transition function $\tau\colon X \to \mathcal{D}_{\mathbb{Q}}(X)$; and
- a labelling function $l\colon X \to L$

◀

Example 6.2.2: Let $X = \{x, x_1, x_2\}$ and $L = \{a, b\}$. We represent the LMC (X, L, τ, l) with $\tau(x) = \frac{1}{2}\,|x_1\rangle + \frac{1}{2}\,|x_2\rangle, \tau(x_1) = 1\,|x_1\rangle, \tau(x_2) = 1\,|x_2\rangle$ and $l(x) = l(x_1) = a, l(x_2) = b$ as:

$$\begin{tikzcd} x_1; a & x; a \arrow[l, "\frac{1}{2}"'] \arrow[r, "\frac{1}{2}"] & x_2; b \end{tikzcd}$$

We use the notation $x; a$ for a state $x \in X$ such that $l(x) = a$. Further, any state with no outgoing edges is assumed to have a self-loop with probability 1. ◀

We now recall a definition of the behavioural distance (henceforth written bd) of states in an LMC as the least fixed point of a functional based on non-expansive maps. This distinguishes two cases: states having different labels should be maximally far apart, so they have distance 1; the distance of states with the same label is then defined recursively, and can be seen as an optimisation problem. Intuition for the distance is most often given in terms of its dual based on couplings under the Kantorovich-Rubinstein duality. The distance between distributions can in that setting be seen as the minimal cost of transporting one unit of mass from one distribution to the other, with the cost of transporting m probability mass between states at distance d being $m \cdot d$. In a more general form (as discussed in [Vil+08]) the distance below can be seen as the maximisation of profit, where we think of buying from one distribution and selling to the other, with the maps h representing buying/selling costs.

On small examples, these intuitions can often be used to determine the distance by examination; on larger systems a solver for linear programs is usually necessary. For this, the definition below can be transformed into a rational linear program, as we use in the proof of Lemma 6.3.8. For further discussions of these distances, see for example [Bre17; RB23]

Definition 6.2.3: For (X, L, τ, l) an LMC, and PMet_X the set of pseudometric spaces on X, we define $\Gamma\colon \mathsf{PMet}_X \to \mathsf{PMet}_X$ by

$$\Gamma(d)(x, y) = \begin{cases} 1, & \text{if } l(x) \neq l(y), \\ \sup_{h\colon (X,d) \to ([0,1], d_e)} \tau(x) \cdot h - \tau(y) \cdot h, & \text{o.w.} \end{cases}$$

Then we define $\mathrm{bd} := \mathrm{lfp}(\Gamma)$. ◂

Note that the least fixed point exists, because PMet_X is a complete lattice, and Γ preserves the pointwise order on PMet_X.

Example 6.2.4: Consider the following LMC:

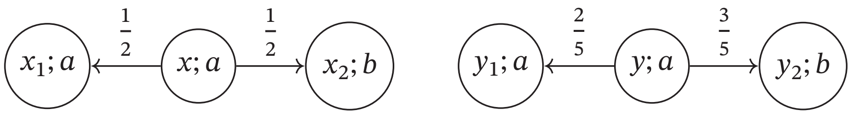

Note that $\mathrm{bd}(x_1, y_1) = \mathrm{bd}(x_2, y_2) = 0$ and $\mathrm{bd}(x_1, y_2) = \mathrm{bd}(x_2, y_1) = 1$, which are the values given by $\Gamma(\bot)$. The value $\mathrm{bd}(x, y)$ is then $\Gamma^2(\bot)(x, y)$ for which it can be shown that the supremum is achieved by the map $h_0(z) =$ **if** $z \in \{x_1, y_1\}$ **then** 1 **else** 0 so that:

$$\Gamma^2(\bot)(x,y) = \sup_{h\colon (X,\Gamma(\bot))\to([0,1],d_e)} \tau(x)\cdot h - \tau(y)\cdot h \tag{6.1}$$

$$= \tau(x)\cdot h_0 - \tau(y)\cdot h_0 \tag{6.2}$$

$$= \left(\frac{1}{2}\cdot 1 + \frac{1}{2}\cdot 0\right) - \left(\frac{2}{5}\cdot 1 + \frac{3}{5}\cdot 0\right) = \frac{1}{10} \tag{6.3}$$

In terms of transportation costs, this $\frac{1}{10}$ can be seen as the cost of transporting the $\frac{1}{10}$ of mass from x_1 to y_2 which can not be transported to y_1 due to its lack of capacity. Intuition based on profits is harder to give in this case, due to the asymmetry of our definition. It is easier to obtain the h in this example by seeing it as grouping the states into “equivalence” classes, spaced as far apart as possible while respecting the distance on X (in this case $\Gamma(\bot)$). ◂

It will be important for the correspondence results of later sections that the behavioural distance can be obtained as a countable supremum, namely the supremum over all finite applications of Γ to the constant zero distance. A similar result for LMCs with non-determinism is shown already in [Des+02, Sec. 3]. It can also be proved using the Kleene fixpoint theorem, or ω-(co)continuity of Γ as shown in [Bre12].

Proposition 6.2.5: For any LMC (X, L, τ, l) and $x, y \in X$, we have

$$\mathrm{bd}(x,y) = \sup_{i<\omega} \Gamma^i(\bot)(x,y)$$ ◂

6.3 Proof System

In this section, we define our derivation system for lower bounds on behavioural distances between states of an LMC. The conclusion of the rules are of the form $x \#_\varepsilon y$ which, as our soundness result will show, implies that $\mathrm{bd}(x, y) \geq \varepsilon$, i.e., the behavioural distance between x and y is at least ε. The definition of bd suggests two rules, one for each case. The label case straightforwardly yields the rule

$$\frac{l(x) \neq l(y)}{x \#_1 y}\ (\text{lab})$$

In the supremum case, $\mathrm{bd}(x, y)$ can be bounded from below by $\tau(x) \cdot h - \tau(y) \cdot h$ for any non-expansive map h by definition. However, it is not immediately clear that this can be done inductively, as we can not assume bd to be known, and thus cannot use it to choose a non-expansive h. Fortunately, as long as the system is sound with respect to the behavioural distance, it suffices to have a map h for which a kind of *pairwise non-expansiveness* holds: for any x', y' we have $|h(x') - h(y')| \leq \varepsilon$ for some ε such that we have proved $x' \#_\varepsilon y'$. Soundness then implies that $|h(x') - h(y')| \leq \varepsilon \leq \mathrm{bd}(x', y')$ for all x', y', which is exactly non-expansiveness of h with respect to bd.

Now, in proofs, we could allow arbitrary recursive proofs and require the choice of a pairwise non-expansive map to correctly apply the rule. Alternatively, we can choose to allow arbitrary maps. We are then required to prove that for all $x', y' \in X$, $|h(x') - h(y')|$ is a lower bound on the behavioural distance. We can see the corresponding proof obligations $x' \#_{|h(x')-h(y')|} y'$ as those generated by a chosen map h. The first option fits with a forward reasoning approach to constructing a proof; we prove some bounds and try to find a fitting h. The second is a backward approach; if we wish to show a bound $x \#_\varepsilon y$, we must supply an h and recursively prove its validity.

We choose the latter approach, primarily because it makes the proof obligations clearer, and we will be able to see when a choice of map is not valid. Using the earlier form, a chosen h may be invalid because we have not proved strong enough bounds, or because it is simply not non-expansive with respect to bd. Such a rule can be written as follows:

$$\frac{h \colon X \to [0,1] \qquad \forall x', y' \in X.\, x' \#_{|h(x')-h(y')|} y' \qquad \tau(x) \cdot h - \tau(y) \cdot h \geq \varepsilon}{x \#_\varepsilon y}$$

However, in this form, the rule does not give a usable derivation system. Namely, in case an LMC has an infinite state space, there will be infinitely many recursive proof obligations. We will ensure that this is always finite, and even reduce the work required to construct proofs beyond this. Further, the map h is so far *real*-valued. This poses a problem both for our approximate completeness result, in which we need to be able to compute these maps, and the construction from proofs to modal formulas whose interpretation will be rational-valued. Summarising, to remedy these issues, we adapt the above rule in the following ways:

- we show that h need only be defined on a finite subset of the state space thereby generating only finitely many proof obligations;
- we restrict the codomain of h to rationals;
- we reduce the number of recursive proof obligations further by not requiring proofs for those bounds which follow from reflexivity and symmetry.

Together, this brings us to the following system:

Definition 6.3.1: Let (X, L, τ, l) be an LMC, $x, y \in X$, and $\varepsilon \in [0,1]_{\mathbb{Q}}$. Further, we define $\mathcal{S}_{x,y} := \{s \in X \mid \tau(x)(s) \neq \tau(y)(s)\}$ and $\mu \cdot_{\mathcal{S}_{x,y}} h := \sum_{s \in \mathcal{S}_{x,y}} \mu(s) \cdot h(s)$.

Then, we define the following derivation rules:

$$\frac{}{x \#_0 y}\,(\text{zero}) \qquad \frac{l(x) \neq l(y)}{x \#_1 y}\,(\text{lab})$$

$$\frac{y \#_\varepsilon x}{x \#_\varepsilon y}\,(\text{symm}) \qquad \frac{x \#_{\varepsilon'} y \qquad \varepsilon \leq \varepsilon'}{x \#_\varepsilon y}\,(\text{weak})$$

$$\frac{\begin{array}{c} h\colon \mathcal{S}_{x,y} \to [0,1]_{\mathbb{Q}} \\ \tau(x) \cdot_{\mathcal{S}_{x,y}} h - \tau(y) \cdot_{\mathcal{S}_{x,y}} h \geq \varepsilon \\ \forall x', y' \in \mathcal{S}_{x,y}.\, h(x') > h(y') \implies x' \#_{h(x') - h(y')} y' \end{array}}{x \#_\varepsilon y}\,(\text{exp})$$

We may drop the subscripts x, y and $\mathcal{S}$ whenever clear from the context. We write $\mathcal{T}_X$ for the set of proof trees inductively defined by these rules, i.e., the smallest set which contains all instances of the *(zero)* and *(lab)* rules for $x, y \in X$, and is closed under applications of all instances of *(symm)*, *(weak)*, and *(exp)* for any $x, y \in X$ and $\varepsilon \in [0,1]_{\mathbb{Q}}$.

Further, we write $\vdash x \#_\varepsilon y$ to mean that the given judgment is provable, i.e., there is a proof tree in $\mathcal{T}_X$ with the given judgment at the root. ◂

Note that we write conditions as premises in the rules rather than as separate side-conditions in the definition for convenience. As suggested by the definition of $\mathcal{T}_X$, this makes both the *(zero)* and *(lab)* rules axioms. In applications of the rules, we may also discuss the side-conditions separately, in order to save space in the proof. Further, strictly speaking, *(exp)* is a family of rules, indexed by the maps h.

As we show in Theorem 6.3.9, the label *(lab)* and expectation *(exp)* rules together already yield a system which is (approximately) complete. For the construction of a proof corresponding to a modal formula (Theorem 6.4.4), we also require the additional weakening *(weak)* rule, essentially to handle formulas using the shift operator ($\varphi \ominus q$, see Definition 6.4.1). To make our proof system and its presentation more pleasant, we include the two further (derivable) zero *(zero)* and symmetry *(symm)* rules. The three additional rules (*(weak)*, *(zero)*, and *(symm)*) are inspired by those from quantitative equational theories [MPP16].

Remark 6.3.2: Note that the restriction to $\mathcal{S}$ in *(exp)* means proofs, which are finite depth by definition, will also be finite breadth even when the LMC under consideration has an infinite state space. This is because $\{s \in X \mid \tau(x)(s) \neq \tau(y)(s)\}$ is a subset of $\mathrm{supp}(\tau(x)) \cup \mathrm{supp}(\tau(y))$, which is finite by assumption. We will see later that this allows us to provide witnesses as both finite proof trees and finite modal formulas. This improves on the earlier work of [RB23] which restricts to finite state spaces. We illustrate this improvement in Example 6.5.6, once we have shown soundness and completeness of the proof system, and its correspondence with modal formulas. ◂

Example 6.3.3: We continue with the LMC from Example 6.2.4 and show how we can prove the distance between x and y shown there as a lower bound. Using the *(lab)* rule, the bounds $u \#_1 v$ can be proved for $u \in \{x_1, y_1\}$ and $v \in \{x_2, y_2\}$. This allows us to define $h_0 : \mathcal{S} \to [0,1]_{\mathbb{Q}}$ as before by $h_0(z) = \textbf{if } z \in \{x_1, y_1\} \textbf{ then } 1 \textbf{ else } 0$ and construct the following proof

$$\dfrac{\dfrac{}{x_1 \#_1 x_2} \qquad \dfrac{}{x_1 \#_1 y_2} \qquad \dfrac{}{y_1 \#_1 x_2} \qquad \dfrac{}{y_1 \#_1 y_2}}{x \#_{\frac{1}{10}} y}$$

which uses $\tau(x) \cdot h_0 - \tau(y) \cdot h_0 = \frac{1}{10}$. ◂

Example 6.3.4: Our second example serves to illustrate a limitation of our proof system, namely that the behavioural distance of states will not always be exactly provable in our system. It would only be provable if we allowed infinite depth proof trees. The LMC we consider is the same as the one in [RB23, Thm. 17], which shows that there is an LMC containing states for which it is not possible to give a single formula "explaining" their distance.

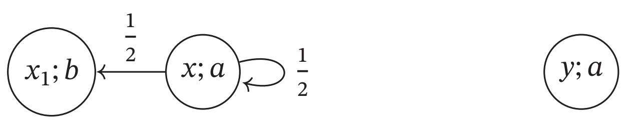

As is discussed in *op. cit.*, the distance $\mathrm{bd}(x, y) = 1$ is reached only in the limit, not by any $\Gamma^i(\bot)$ and thus not by any single tree. Proving the bound given by $\Gamma^i(\bot)$ (for $i > 0$) can be done using $i - 1$ applications of the *(exp)* rule together with two applications each of the *(lab)* and *(zero)* rules. For example, once we have proved $x_1 \#_1 u$ for $u \in \{x, y\}$, we can take the map $h_0 \colon x, y \mapsto 0, x_1 \mapsto 1$ for which $\tau(x) \cdot h_0 - \tau(y) \cdot h_0 = \frac{1}{2}$ and prove:

$$\frac{\dfrac{}{x_1 \#_1 x} \qquad \dfrac{}{x_1 \#_1 y}}{x \#_{\frac{1}{2}} y}$$

This step (plus an application of *(symm)*) allows the next application of *(exp)* with a new non-expansive $h_0 \colon x \mapsto \frac{1}{2}, y \mapsto 0, x_1 \mapsto 1$ (increasing the value of the previous h_0 from 0 to $\frac{1}{2}$), yielding a bound of $\frac{3}{4}$. Continuing to increase the value of $h_0(x)$ in this way, we approach $\mathrm{bd}(x, y)$ from below. ◂

6.3.1 Soundness

We now move on to showing soundness of the system, i.e., that if we can prove $x \#_\varepsilon y$, then $\mathrm{bd}(x, y) \geq \varepsilon$. The *(zero)* rule is sound as our pseudometrics are valued in $[0, 1]$ and thus 0 is always a sound lower bound. Similarly, the behavioural distance is symmetric, so that the order of states does not change a lower bound's validity. Our weakening rule is sound by transitivity of $\leq$. Soundness of the label rule follows from the definition of Γ. This is similar for the expectation rule, as we discussed at the beginning of this section. We make the intuition given there precise in the coming results.

Due to our restriction of the domain of the map in the *(exp)* rule, we will require the following lemma in the soundness proof:

Lemma 6.3.5: For (X, L, τ, l) an LMC, $d \colon X \times X \to [0, 1]$ a pseudometric and $x, y \in X$:

$$\sup_{h \colon (X,d) \to ([0,1],d_e)} \tau(x) \cdot h - \tau(y) \cdot h = \sup_{h \colon (\mathcal{S},d|_\mathcal{S}) \to ([0,1],d_e)} \tau(x) \cdot_\mathcal{S} h - \tau(y) \cdot_\mathcal{S} h$$

where $d|_\mathcal{S} = d \circ (\iota_\mathcal{S} \times \iota_\mathcal{S})$ with $\iota_\mathcal{S} \colon \mathcal{S} \hookrightarrow X$ the inclusion map. ◂

Proof. We prove two inequalities:

$\leq$: This holds because any $h \colon (X, d) \to ([0, 1], d_e)$ restricts to a map $h|_\mathcal{S} = h \circ \iota_\mathcal{S} \colon (\mathcal{S}, d|_\mathcal{S}) \to ([0, 1], d_e)$, and we can show that

$$\tau(x) \cdot h - \tau(y) \cdot h = \tau(x) \cdot_\mathcal{S} h|_\mathcal{S} - \tau(y) \cdot_\mathcal{S} h|_\mathcal{S}$$

$\geq$: For this direction, we show that for any $h\colon (\mathcal{S}, d|_{\mathcal{S}}) \to ([0,1], d_e)$, there is an $h'\colon (X, d) \to ([0,1], d_e)$ such that

$$\tau(x) \cdot h' - \tau(y) \cdot h' = \tau(x) \cdot_{\mathcal{S}} h - \tau(y) \cdot_{\mathcal{S}} h$$

We use an existing construction of extensions of non-expansive maps, to extend h along the inclusion $\iota_{\mathcal{S}}\colon \mathcal{S} \hookrightarrow X$. Namely, we define $h'(x) := \inf_{z \in \mathcal{S}} h(z) \oplus d(x,z)$, where $\oplus$ is truncated addition on the unit interval. This is an extension in the sense that $h' \circ \iota_{\mathcal{S}} = h$, so that the above equality indeed holds. ■

Remark 6.3.6: We can further restrict the space of possible maps h in the above, by seeing the supremum as the optimal solution of a linear program as follows. We encode functions $h\colon \mathcal{S} \to [0,1]$ as (finite) vectors $\vec{h} \in [0,1]^{|\mathcal{S}|}$, writing $\vec{h}_x$ for the value of $\vec{h}$ at the position indexed by $x \in \mathcal{S}$. Then, each inequality $|h(x) - h(y)| \leq d(x,y)$ can be expressed by $\vec{a} \cdot \vec{h} \leq d(x,y)$ and $\vec{a}' \cdot \vec{h} \leq d(x,y)$ with $\vec{a}_x = 1, \vec{a}_y = -1, \vec{a}'_x = -1, \vec{a}'_y = 1$ (and all other entries zero). We can enforce $0 \leq \vec{h}_x \leq 1$ similarly. We are thus interested in the problem of maximising $\tau(x) \cdot h - \tau(y) \cdot h$ (a linear expression) subject to the constraints expressed by $\mathbf{A} \cdot \vec{h} \leq \vec{b}$ for an integer matrix $\mathbf{A}$ and vector $\vec{b}$.

The feasible region of this problem is a polytope, which we call $H^d_{x,y}$. It is convex and closed due to the shape of the constraints, and bounded due to the restriction to the unit interval. It is furthermore non-empty as any constant map valued in $[0,1]$ lies within it. It is known (see, e.g., [Tru71]) that for such a feasible region, the optimal value is achieved in the vertices of the polytope, which we denote by $V(H^d_{x,y})$. Then we can write

$$\sup_{h\colon (X,d) \to ([0,1], d_e)} \tau(x) \cdot h - \tau(y) \cdot h = \max_{h \in V(H^d_{x,y})} \tau(x) \cdot_{\mathcal{S}} h - \tau(y) \cdot_{\mathcal{S}} h$$

The supremum becomes a maximum because finitely many linear inequalities define a polytope with finitely many vertices.

Obtaining a map h in which the supremum is achieved can thus be done using an algorithm such as simplex. We will apply this recursively for $d = \Gamma^i(\bot)$ (the finite approximants of bd), to construct proof trees in our proof of approximate completeness in the next section. ◂

We are now able to show (by structural induction) that any proof in our system yields a lower bound on bd.

Theorem 6.3.7 (Soundness): For any LMC (X, L, τ, l), any proof tree built from the rules of Definition 6.3.1, any $\varepsilon \in [0,1]_{\mathbb{Q}}$, and any $x, y \in X$, if the proof tree has $x \#_\varepsilon y$ at the root, then $\mathrm{bd}(x,y) \geq \varepsilon$. ◂

Proof. We proceed by induction on the structure of the proof tree.

Case *(zero)*: By its definition, bd takes values in $[0,1]$, so that $\mathrm{bd}(x,y) \geq 0$ holds.

Case *(label)*: We have a proof tree

$$\frac{l(x) \neq l(y)}{x \#_1 y}$$

By definition of Γ, we must have $\mathrm{bd}(x,y) = 1$, so that indeed $\mathrm{bd}(x,y) \geq 1$.

Case *(symm)*: We have a proof tree

$$\frac{y \#_\varepsilon x}{x \#_\varepsilon y}$$

By induction, we have $\mathrm{bd}(y,x) \geq \varepsilon$, but bd is symmetric, so that also $\mathrm{bd}(x,y) \geq \varepsilon$.

Case *(weak)*: We have a proof tree

$$\frac{x \#_{\varepsilon'} y \qquad \varepsilon \leq \varepsilon'}{x \#_\varepsilon y}$$

By induction, we have $\mathrm{bd}(x,y) \geq \varepsilon' \geq \varepsilon$.

Case *(exp)*: We have a proof tree

$$\frac{\begin{array}{c} h\colon \mathcal{S} \to [0,1]_{\mathbb{Q}} \\ \tau(x)\cdot h - \tau(y)\cdot h \geq \varepsilon \\ \forall x', y' \in \mathcal{S}.\, h(x') > h(y') \implies x' \#_{h(x')-h(y')} y' \end{array}}{x \#_\varepsilon y}\ (\mathrm{exp})$$

By induction, we have for all $x', y' \in \mathcal{S}$ with $h(x') > h(y')$ that $\mathrm{bd}(x',y') \geq |h(x') - h(y')|$. For $x', y' \in \mathcal{S}$ with $h(x') < h(y')$, we have $\mathrm{bd}(x',y') = \mathrm{bd}(y',x') \geq |h(y') - h(x')| = |h(x') - h(y')|$. For the remaining pairs, $|h(x') - h(y')| = 0 \leq \mathrm{bd}(x',y')$. In other words, $h\colon \mathcal{S} \to [0,1]_{\mathbb{Q}}$ is a non-expansive map $h\colon (\mathcal{S}, \mathrm{bd}) \to ([0,1]_{\mathbb{Q}}, d_e)$. As bd is defined as the least fixed point of Γ, we have

$$\begin{aligned} \mathrm{bd}(x,y) &= \Gamma(\mathrm{bd})(x,y) \\ &= \begin{cases} 1, & \text{if } l(x) \neq l(y), \\ \sup_{h\colon (X,\mathrm{bd}) \to ([0,1], d_e)} \tau(x)\cdot h - \tau(y)\cdot h, & \text{o.w.} \end{cases} \end{aligned}$$

In case $l(x) \neq l(y)$, we have $\mathrm{bd}(x,y) = 1 \geq \varepsilon$.

In the remaining case, we have

$$\mathrm{bd}(x,y) = \sup_{k\colon (X,\mathrm{bd})\to([0,1],d_e)} \tau(x)\cdot k - \tau(y)\cdot k \tag{6.4}$$

$$= \sup_{h\colon (\mathcal{S},d|_{\mathcal{S}})\to([0,1],d_e)} \tau(x)\cdot_{\mathcal{S}} h - \tau(y)\cdot_{\mathcal{S}} h \tag{6.5}$$

$$\geq \tau(x)\cdot h - \tau(y)\cdot h \geq \varepsilon \tag{6.6}$$

where the second equality is shown in Lemma 6.3.5 and the first inequality holds because h is one of the non-expansive maps ranged over in the sup. This covers all cases, so soundness follows by induction. ∎

6.3.2 Approximate Completeness

The rest of this section is dedicated to proving the approximate completeness of the system. This will show that we can prove lower bounds arbitrarily close to the "true" value given by bd. Our proof relies on the fact that we can get arbitrarily close to bd with its finite approximants $\Gamma^i(\bot)$, and the following lemma, that shows how we can construct proofs exhibiting these finite approximants as lower bounds on the behavioural distance.

Lemma 6.3.8: For any $i \in \mathbb{N}$ and $x, y \in X$, we have $\vdash x \#_{\Gamma^i(\bot)(x,y)} y$. ◂

Proof (by induction on i). We strengthen the induction by showing that the $\Gamma^i(\bot)(x,y)$ are always rational. For the base case, we have $\Gamma^0(\bot)(x,y) = \bot(x,y) = 0$ which is rational, and we can prove $x \#_0 y$ using the *(zero)* rule.

Now let $i \in \mathbb{N}$, and suppose for any $x, y \in X$ that:

1. we can prove $x \#_{\Gamma^i(\bot)(x,y)} y$
2. $\Gamma^i(\bot)(x,y)$ is rational.

We have

$$\begin{aligned}
&\Gamma^{i+1}(\bot)(x,y)\\
&= \Gamma(\Gamma^i(\bot))(x,y)\\
&= \begin{cases} 1, & \text{if } l(x) \neq l(y),\\ \sup_{h\colon (X,\Gamma^i(\bot))\to([0,1],d_e)} \tau(x)\cdot h - \tau(y)\cdot h, & \text{o.w.} \end{cases}
\end{aligned}$$

If $l(x) \neq l(y)$, we can prove $x \#_1 y$ using the *(lab)* rule. Otherwise, by Lemma 6.3.5, we have

$$\sup_{h\colon (X,\Gamma^i(\bot))\to([0,1],d_e)} \tau(x)\cdot h - \tau(y)\cdot h \tag{6.7}$$

$$= \sup_{h\colon (\mathcal{S},\Gamma^i(\bot)|_{\mathcal{S}})\to([0,1],d_e)} \tau(x)\cdot_{\mathcal{S}} h - \tau(y)\cdot_{\mathcal{S}} h \tag{6.8}$$

As explained in Remark 6.3.6, we can find an optimal h_0 by solving a linear program. This will in particular be a map $h_0 : (\mathcal{S}, \Gamma^i(\bot)|_{\mathcal{S}}) \to ([0,1]_{\mathbb{Q}}, d_e)$ (note the restriction to rationals) because all the coefficients in the problem are rational, by induction. We thus have

$$\sup_{h : (\mathcal{S}, \Gamma^i(\bot)|_{\mathcal{S}}) \to ([0,1], d_e)} \tau(x) \cdot_{\mathcal{S}} h - \tau(y) \cdot_{\mathcal{S}} h = \tau(x) \cdot_{\mathcal{S}} h_0 - \tau(y) \cdot_{\mathcal{S}} h_0$$

which is rational, because the transition probabilities given by $\tau(x)$ and $\tau(y)$ are also rational. We can now construct the following proof

$$\frac{\begin{array}{c} h_0 : \mathcal{S} \to [0,1]_{\mathbb{Q}} \\ \tau(x) \cdot h_0 - \tau(y) \cdot h_0 \geq \varepsilon \\ \forall x', y' \in \mathcal{S}.\, h_0(x') > h_0(y') \implies x' \#_{h_0(x') - h_0(y')} y' \end{array}}{x \#_\varepsilon y}$$

in which recursive proofs are given by induction, and the above equality together with the definition of the approximants $\Gamma^n(\bot)$ allows us to choose $\varepsilon := \Gamma^{i+1}(\bot)(x, y)$. ■

We see that completeness in this sense only requires the use of the *(zero)*, *(lab)*, and *(exp)* rules. In fact, it can be shown for only the latter two, with the *(zero)* rule being essentially an instance of the *(exp)* rule where the map h is taken to be constant. The approximate completeness of the system is now a simple consequence of Lemma 6.3.8 and Proposition 6.2.5, with no further applications of the rules needed.

Theorem 6.3.9 (Approximate Completeness): For any (real) $\delta > 0$, and any $x, y \in X$ there is a proof tree with $x \#_\varepsilon y$ at the root, so that $0 \leq \mathrm{bd}(x, y) - \varepsilon < \delta$. ◂

Proof. Let $\delta > 0$. By Proposition 6.2.5, we have $\mathrm{bd}(x, y) = \sup_{i<\omega} \Gamma^i(\bot)(x, y)$, so there exists $i \in \mathbb{N}$ such that

$$0 \leq \mathrm{bd}(x, y) - \Gamma^i(\bot)(x, y) < \delta$$

Lemma 6.3.8 exactly gives us a proof of $x \#_{\Gamma^i(\bot)(x,y)} y$, and we are done. ■

6.4 Logic

The previous sections have given us a way to inductively derive lower bounds on the behavioural distance between states of an LMC, and shown soundness and approximate completeness of the proof system with respect to the behavioural distance bd. In this sense, the system gives finite evidence or *witnesses* for (lower bounds on) behavioural distances.

Another approach to giving such evidence is to construct formulas in some logic which, in the terminology of [RB23], “explain” the difference. Ideally, given states $x, y \in X$, this would be a formula φ such that $|[\![\varphi]\!](x) - [\![\varphi]\!](y)| = \mathrm{bd}(x, y)$, i.e., the difference in interpretations of φ on the states is exactly equal to their behavioural distance. However, as is shown in *op. cit.*, such a formula can not be given in general (cf. Example 6.3.4). Instead, a construction is given of formulas corresponding to finite approximations of bd. In our notation, these are formulas φ such that $|[\![\varphi]\!](x) - [\![\varphi]\!](y)| = \Gamma^i(\bot)(x, y)$ (for some $i \in \mathbb{N}$). As discussed in Section 6.2, bd can be obtained as the countable limit of these approximations, so that the construction of [RB23] gives formulas explaining the behavioural distance of states up to an arbitrarily small error.

In this section, we give analogous constructions between proofs and formulas in the same logic used to characterise the behavioural distance bd in [RB23]. We start by recalling this logic and its interpretation on LMCs. Its semantics is given in terms of a real-valued interpretation function (first suggested in [Koz85]) and is a slight variation of the logic studied in relation to behavioural distances in [Des+04]. We then move onto the related constructions, the first of which is a straightforward inductive construction of a proof that the distance $|[\![\varphi]\!](x) - [\![\varphi]\!](y)|$ is a lower bound. The second, again inductively, constructs a formula witnessing some proved lower bound. This is based on constructions in [RB23] and relies on a non-trivial lemma in the case where the distance of states arises from the supremum case in the definition of Γ.

Definition 6.4.1: Define the syntax of the logic $\mathcal{L}$ by the following grammar:

$$\varphi ::= a \mid \bigcirc\varphi \mid \neg\varphi \mid \varphi \ominus q \mid \varphi \vee \varphi \tag{6.9}$$

where $a \in L$ and $q \in [0, 1]_{\mathbb{Q}}$. Further, given an LMC (X, L, τ, l), the quantitative semantics of $\mathcal{L}$ is given by the equations

$$[\![a]\!](x) = \begin{cases} 1, & \text{if } l(x) = a, \\ 0, & \text{o.w.} \end{cases} \qquad [\![\varphi \ominus q]\!](x) = \max(0, [\![\varphi]\!](x) - q)$$

$$[\![\bigcirc\varphi]\!](x) = \tau(x) \cdot [\![\varphi]\!] \qquad [\![\varphi \vee \psi]\!](x) = \max([\![\varphi]\!](x), [\![\psi]\!](x))$$

$$[\![\neg\varphi]\!](x) = 1 - [\![\varphi]\!](x)$$

which define the interpretation function $[\![\cdot]\!] : \mathcal{L} \to X \to [0, 1]_{\mathbb{Q}}$ recursively. ◀

Remark 6.4.2: From the connectives in the logic $\mathcal{L}$, it is possible to define also $\wedge$ and $\oplus$ as $\varphi \wedge \psi := \neg(\neg\varphi \vee \neg\psi)$ and $\varphi \oplus q := \neg(\neg\varphi \ominus q)$, which then have the expected semantics. We also write false for the formula $a \ominus 1$ whose interpretation is everywhere zero. ◂

Example 6.4.3: Consider the LMC from Example 6.3.4, and the formulas $\varphi_i := \bigcirc^i b$ for $i \in \mathbb{N}$. We can show that $[\![\varphi_i]\!](x_1) = 1$ for any i, so that $[\![\varphi_i]\!](x) = \sum_{n=1}^{i} \left(\frac{1}{2}\right)^i$, while $[\![\varphi_i]\!](y) = 0$. The formula φ_i captures the probability of reaching a state with label b after i steps. ◂

The logic and its interpretation induce new distances between states, namely the difference in interpretations $|[\![\varphi]\!](x) - [\![\varphi]\!](y)|$. Our first correspondence result shows that this distance can be shown to be a lower bound on $\mathrm{bd}(x, y)$ by a proof in our system.

Theorem 6.4.4: For any LMC (X, L, τ, l), formula $\varphi \in \mathcal{L}$, and $x, y \in X$, there exists a proof tree with $x \#_\varepsilon y$ as its root, where $\varepsilon = |[\![\varphi]\!](x) - [\![\varphi]\!](y)|$. ◂

Proof (by induction on the structure of φ). We write $\mathcal{T}(\psi, x, y)$ for the proof tree constructed for a formula ψ and states x, y.

Case $\varphi = a$: We have $\varepsilon = 0$ or $\varepsilon = 1$ with the following proofs:

$$\frac{}{x \#_0 y}\ (\text{zero}) \qquad \frac{l(x) \neq l(y)}{x \#_1 y}\ (\text{lab})$$

Case $\varphi = \neg\psi$: We have $|[\![\varphi]\!](x) - [\![\varphi]\!](y)| = |[\![\psi]\!](x) - [\![\psi]\!](y)|$ so simply let $\mathcal{T}(\varphi, x, y) = \mathcal{T}(\psi, x, y)$.

Case $\varphi = \psi \ominus q$: In this case we will have $|[\![\psi]\!](x) - [\![\psi]\!](y)| \leq |[\![\varphi]\!](x) - [\![\varphi]\!](y)|$ (truncation may give an inequality) so that we have

$$\frac{x \#_{|[\![\varphi]\!](x) - [\![\varphi]\!](y)|}\, y \qquad |[\![\psi]\!](x) - [\![\psi]\!](y)| \leq |[\![\varphi]\!](x) - [\![\varphi]\!](y)|}{x \#_{|[\![\psi]\!](x) - [\![\psi]\!](y)|}\, y}\ (\text{weak})$$

Case $\varphi = \varphi_1 \vee \varphi_2$: We have

$$\varepsilon = |[\![\varphi_1 \vee \varphi_2]\!](x) - [\![\varphi_1 \vee \varphi_2]\!](y)| \tag{6.10}$$
$$= |\max([\![\varphi_1]\!](x), [\![\varphi_2]\!](x)) - \max([\![\varphi_1]\!](y), [\![\varphi_2]\!](y))| \tag{6.11}$$
$$\leq \max(|[\![\varphi_1]\!](x) - [\![\varphi_1]\!](y)|, |[\![\varphi_2]\!](x) - [\![\varphi_2]\!](y)|) \tag{6.12}$$

so that we take $\mathcal{T}(\varphi, x, y)$ to be

$$\frac{\mathcal{T}(\varphi_i, x, y) \qquad \varepsilon \leq \varepsilon'}{x \#_\varepsilon y}\ (\text{weak})$$

with φ_i the formula yielding the above maximum, which we have called ε'.

Case $\varphi = \bigcirc\psi$: Here, we have $\varepsilon = |\tau(x)\cdot[\![\psi]\!]-\tau(y)\cdot[\![\psi]\!]|$ and by induction, we have for all $x', y' \in \mathcal{S}$ with $[\![\psi]\!](x') > [\![\psi]\!](y')$ trees $\mathcal{T}(\psi, x', y')$. Now we must distinguish two cases. If $\tau(x)\cdot[\![\psi]\!] \geq \tau(y)\cdot[\![\psi]\!]$, we take $h = [\![\psi]\!]|_{\mathcal{S}}$ and construct

$$\frac{\begin{array}{l}[\![\psi]\!]|_{\mathcal{S}} : \mathcal{S} \to [0,1]_{\mathbb{Q}} \\ \varepsilon = \tau(x)\cdot[\![\psi]\!] - \tau(y)\cdot[\![\psi]\!] \qquad \{\mathcal{T}(\psi, x', y') \mid h(x') > h(y')\}\end{array}}{x \#_\varepsilon y}\ (\text{exp})$$

Otherwise, we take $h = [\![\neg\psi]\!]|_{\mathcal{S}}$, and have

$$\frac{\begin{array}{l}[\![\neg\psi]\!]|_{\mathcal{S}} : \mathcal{S} \to [0,1]_{\mathbb{Q}} \\ \varepsilon = \tau(x)\cdot[\![\neg\psi]\!] - \tau(y)\cdot[\![\neg\psi]\!] \qquad \{\mathcal{T}(\psi, x', y') \mid h(x') > h(y')\}\end{array}}{x \#_\varepsilon y}\ (\text{exp})$$

Note that

$$\tau(x)\cdot[\![\neg\psi]\!] - \tau(y)\cdot[\![\neg\psi]\!] = (1-\tau(x)\cdot[\![\psi]\!]) - (1-\tau(y)\cdot[\![\psi]\!]) \tag{6.13}$$
$$= \tau(y)\cdot[\![\psi]\!] - \tau(x)\cdot[\![\psi]\!] \tag{6.14}$$

Under the assumption that $\tau(y)\cdot[\![\psi]\!] \geq \tau(x)\cdot[\![\psi]\!]$, this equals ε. This completes the construction. ■

Remark 6.4.5: The depth of the constructed proof matches the modal depth of the formula, i.e., the maximum number of nested $\bigcirc$ modalities. In this sense, the proof does not grow unexpectedly compared to the formula we start with. Due to the branching in the *(exp)* rule, the number of rules we apply will be larger than the number of operators in a given formula in general. For an example, recall the LMC from Example 6.2.4, and consider the formula $\bigcirc a$. This has just two operators, but the proof generated by the above procedure will contain five recursive proof obligations generated by $[\![\bigcirc a]\!]$ in the application of *(exp)*, each requiring at least one rule application. ◂

6.5 Constructing formulas from proofs

To complete the correspondence between proofs and modal formulas, in this section we provide a construction going from a proof of $x \#_\varepsilon y$, to a formula φ such that $|[\![\varphi]\!](x) - [\![\varphi]\!](y)| = \varepsilon$. In fact, we construct formulas whose interpretation on x is equal to the lower bound, and whose interpretation on y is zero. This is required for our recursive construction.

Our construction is inspired by that of [RB23], which relies on the following lemma:

Lemma 6.5.1: Let $f : X \to [0,1]$. If for any $x, y \in X$, we have a function $g_{xy} : X \to [0,1]$ such that $g_{xy}(x) = f(x)$ and $g_{xy}(y) = f(y)$ then $f = \max_x \min_y g_{xy} = \min_x \max_y g_{xy}$. ◂

A proof of a more general (continuous) version of this result can be found in [Ash72, Lemma A7.2], where it is used in the proof of the Stone-Weierstrass Theorem.

The usefulness of this lemma is perhaps not very intuitive; the idea is that given an application of the *(exp)* rule, we may recursively assume formulas $\varphi_{x'y'}$ capturing the distance between successor states x', y', i.e., $\varphi_{x'y'}(x') = h(x') - h(y')$ and $\varphi_{x'y'}(y') = 0$. We thus have formulas whose interpretations match the function h, but only for pairs of states, so that a lemma similar to the above is required to combine them into a single formula.

Note that, due to the form of our *(exp)* rule, we can not assume in an inductive proof that formulas are given for all successors; only those x', y' for which $h(x') > h(y')$. It is possible to recover all other pairs using the *(zero)* and *(symm)* rules, but we wish to keep the formulas as small as we can. For this, we prove the following stronger version of Lemma 6.5.1.

Lemma 6.5.2: Let $f : X \to [0,1]$. If for any $x, y \in X$ such that $f(x) \geq f(y)$, we have a function $g_{xy} : X \to [0,1]$ such that:

(1) $g_{xy}(x) = f(x)$,

(2) $g_{xy}(y) = f(y)$,

(3) $\forall z \in X.\, g_{xy}(z) \geq f(y)$,

(4) $\forall z \in X.\, g_{xx}(z) = f(x)$,

then $f = \max_x \min_{y : f(x) \geq f(y)} g_{xy}$. ◂

Proof. We first define

$$k_{xy} = \begin{cases} g_{xy} & \text{if } f(x) \geq f(y) \\ g_{yx} & \text{if } f(x) < f(y) \end{cases} \tag{6.15}$$

Note that these k_{xy} satisfy the conditions of Lemma 6.5.1 so that $f = \max_x \min_y k_{xy}$. It thus suffices to prove that

$$\max_x \min_y k_{xy} = \max_x \min_{y : f(x) \geq f(y)} g_{xy} \tag{6.16}$$

We prove this by proving two inequalities.

$\leq$: Consider the inequality and simplify as follows:

$$\begin{aligned}
&\forall z \in X.\, \max_x \min_y k_{xy}(z) \leq \max_x \min_{y:\, f(x) \geq f(y)} g_{xy}(z) \\
\iff &\forall z \in X.\, \forall u_1 \in X.\, \min_y k_{u_1 y}(z) \leq \max_x \min_{y:\, f(x) \geq f(y)} g_{xy}(z) \\
\iff &\forall z \in X.\, \forall u_1 \in X.\, \exists u_3 \in X.\, \min_y k_{u_1 y}(z) \leq \min_{y:\, f(u_3) \geq f(y)} g_{u_3 y}(z) \\
\iff &\forall z \in X.\, \forall u_1 \in X.\, \exists u_3 \in X.\, \forall u_4 \in X.\, [f(u_3) \geq f(u_4) \\
&\qquad\qquad \implies \min_y k_{u_1 y}(z) \leq g_{u_3 u_4}(z)] \\
\iff &\forall z \in X.\, \forall u_1 \in X.\, \exists u_3 \in X.\, \forall u_4 \in X.\, [f(u_3) \geq f(u_4) \\
&\qquad\qquad \implies \exists u_2 \in X.\, k_{u_1 u_2}(z) \leq g_{u_3 u_4}(z)]
\end{aligned}$$

So, let $z, u_1 \in X$ and take $u_3 = z$. Further, let $u_4 \in X$ such that $f(u_3) \geq f(u_4)$ and take $u_2 = z$. Then

$$\begin{aligned}
k_{u_1 u_2}(z) &= k_{u_1 z}(z) = f(z) \\
g_{u_3 u_4}(z) &= g_{z u_4}(z) = f(z)
\end{aligned}$$

This first inequality thus holds.

$\geq$: Consider the inequality and simplify as follows:

$$\begin{aligned}
&\forall z \in X.\, \max_x \min_y k_{xy}(z) \geq \max_x \min_{y: f(x) \geq f(y)} g_{xy}(z) \\
\iff &\forall z \in X.\, \forall u_3 \in X.\, \max_x \min_y k_{xy}(z) \geq \min_{y: f(u_3) \geq f(y)} g_{u_3 y}(z) \\
\iff &\forall z \in X.\, \forall u_3 \in X.\, \exists u_1 \in X.\, \min_y k_{u_1 y}(z) \geq \min_{y: f(u_3) \geq f(y)} g_{u_3 y}(z) \\
\iff &\forall z \in X.\, \forall u_3 \in X.\, \exists u_1 \in X.\, \forall u_2 \in X.\, k_{u_1 u_2}(z) \geq \min_{y: f(u_3) \geq f(y)} g_{u_3 y}(z) \\
\iff &\forall z \in X.\, \forall u_3 \in X.\, \exists u_1 \in X.\, \forall u_2 \in X.\, \exists u_4 \in X.\, f(u_3) \geq f(u_4) \land \\
&\qquad\qquad k_{u_1 u_2}(z) \geq g_{u_3 u_4}(z).
\end{aligned}$$

So, we let $z, u_3 \in X$ and take $u_1 = u_3$. Now let $u_2 \in X$. We distinguish two further cases:

$f(u_2) > f(u_3)$: Here we take $u_4 = u_3$ and have

$$k_{u_1 u_2}(z) = k_{u_3 u_2}(z) = g_{u_2 u_3}(z) \overset{(3)}{\geq} f(u_3) \overset{(4)}{=} g_{u_3 u_3}(z) = g_{u_3 u_4}(z)$$

$f(u_2) \leq f(u_3)$: Here we take $u_4 = u_2$ and see that

$$k_{u_1u_2}(z) = k_{u_3u_2}(z) = g_{u_3u_2}(z) = g_{u_3u_4}(z)$$

This concludes the case distinctions. ■

The proof is a little involved and not very informative, however it allows us to construct (potentially) smaller formulas witnessing distances as compared to [RB23]. This is because we do not require formulas distinguishing all pairs of reachable states in the construction of the next theorem. We discuss the improvement in size at the end of this section, and illustrate it in Examples 6.5.5 and 6.5.6.

Theorem 6.5.3: For any LMC (X, L, τ, l), any $\varepsilon \in [0,1]_{\mathbb{Q}}$, and any $x, y \in X$, if we have a proof of $x \#_\varepsilon y$ using the rules of Definition 6.3.1, then there is a formula $\varphi_{xy} \in \mathcal{L}$ such that $[\![\varphi_{xy}]\!](x) = \varepsilon$ and $[\![\varphi_{xy}]\!](y) = 0$. ◂

Note that we abuse notation, and use $x \#_\varepsilon y$ to refer to both a judgment in a proof, and a proof tree with this judgment at the root.

Proof. This is by induction on the structure of the proof.

Case *(zero)*: In this case we take $\varphi_{xy} = \mathsf{false}$. We have $[\![\varphi_{xy}]\!](z) = 0$ for any $z \in X$, yielding the desired interpretations.

Case *(lab)*: Here, we take $\varphi_{xy} = l(x)$. We must have $[\![l(x)]\!](y) = 0$ as $l(x) \neq l(y)$, so that the interpretations are as required.

Case *(symm)*: The induction hypothesis gives a formula φ_{yx}. Taking $\varphi_{xy} = \neg\varphi_{yx} \ominus (1-\varepsilon)$, we have $[\![\varphi_{xy}]\!](x) = (1 - [\![\varphi_{yx}]\!](x)) \ominus (1-\varepsilon) = \varepsilon$ and $[\![\varphi_{xy}]\!](y) = (1 - [\![\varphi_{yx}]\!](y)) \ominus (1-\varepsilon) = 0$ as desired.

Case *(weak)*: The induction hypothesis here gives φ'_{xy} with $[\![\varphi'_{xy}]\!](x) = \varepsilon'$ and $[\![\varphi'_{xy}]\!](y) = 0$ and $\varepsilon' \geq \varepsilon$. This means we can take $\varphi_{xy} = \varphi'_{xy} \ominus (\varepsilon' - \varepsilon)$ and have $[\![\varphi_{xy}]\!](x) = \varepsilon' \ominus (\varepsilon' - \varepsilon) = \varepsilon$ and $[\![\varphi_{xy}]\!](y) = 0 - (\varepsilon' - \varepsilon) = 0$.

Case *(exp)*: Suppose we have a proof of the form

$$\frac{\begin{array}{c} h : \mathcal{S} \to [0,1]_{\mathbb{Q}} \\ \tau(x) \cdot h - \tau(y) \cdot h \geq \varepsilon \\ \forall x', y' \in \mathcal{S}.\, h(x') > h(y') \implies x' \#_{h(x')-h(y')} y' \end{array}}{x \#_\varepsilon y}$$

By induction, we have formulas $\varphi_{x'y'}$ such that $[\![\varphi_{x'y'}]\!](x') = h(x') - h(y')$ and $[\![\varphi_{x'y'}]\!](y') = 0$ for all $x', y' \in \mathcal{S}$ with $h(x') > h(y')$. For those x', y' such that $h(x') = h(y')$ we define $\varphi_{x'y'} := \mathsf{false}$ (any formula which is everywhere zero can be used). Using these we construct, for x', y' such that $h(x') \geq h(y')$, the formulas $\psi^h_{x'y'} := \varphi_{x'y'} \oplus h(y')$.

We now claim that the interpretations $[\![\psi^h_{x'y'}]\!]$ satisfy the conditions of Lemma 6.5.2, so that the construction there allows us to define a formula whose interpretation is exactly h. The first two conditions clearly hold. For the third, note that $[\![\varphi_{x'y'}]\!](z) \geq 0$ for any z, so that indeed

$$[\![\psi^h_{x'y'}]\!](z) = [\![\varphi_{x'y'}]\!](z) \oplus h(y') \geq 0 \oplus h(y') = h(y') \tag{6.17}$$

For the fourth, we see that for any $z \in X$:

$$[\![\psi^h_{x'x'}]\!](z) = [\![\text{false} \oplus h(x')]\!](z) = h(x') \tag{6.18}$$

We now define

$$\varphi^h_{xy} := \bigvee_{x'} \left[\left[\bigwedge_{y' : h(x') > h(y')} \psi^h_{x'y'} \right] \wedge (\text{false} \oplus h(x')) \right] \tag{6.19}$$

This has the same interpretation as $\bigvee_{x'} \bigwedge_{y' : h(x') \geq h(y')} \psi^h_{x'y'}$, because for pairs (x', y') with $h(x') = h(y')$, the formula $\psi^h_{x'y'}$ will be equal to false $\oplus$ $h(x')$ by definition. Note also that these formulas are finite, as we quantify over $\mathcal{S}$. Thus letting $\varphi_{xy} := \bigcirc\varphi^h_{xy} \ominus (\tau(y) \cdot h)$, yields a (finitary) formula with the desired property. ∎

One may wonder why we have constructed the formulas φ^h_{xy} as a disjunction over a conjunction, and not vice versa. It turns out that this order matters in the case of our *(exp)* rule, as the following example shows.

Example 6.5.4: Consider the following LMC:

$$x_1; a \xleftarrow{\frac{1}{2}} x_0; a \xrightarrow{\frac{1}{2}} x_2; b \qquad\qquad y_1; c \xleftarrow{\frac{1}{2}} y_0; a \xrightarrow{\frac{1}{2}} y_2; b$$

For this example, we can prove the bound $y_0 \#_{\frac{1}{2}} x_0$ using the *(exp)* rule and the map $h(x) = \textbf{if } x = x_1 \textbf{ then } 0 \textbf{ else } 1$:

$$\frac{\dfrac{l(x_2) \neq l(x_1)}{x_2 \#_1 x_1} \qquad \dfrac{l(y_1) \neq l(x_1)}{y_1 \#_1 x_1} \qquad \dfrac{l(y_2) \neq l(x_1)}{y_2 \#_1 x_1}}{x_0 \#_{\frac{1}{2}} y_0} \tag{6.20}$$

Constructing the $\psi^h_{x'y'}$ for only the pairs occurring in recursive proofs yields

$$\psi^h_{x_2x_1} = b \oplus 0 \quad \psi^h_{y_1x_1} = c \oplus 0 \quad \psi^h_{y_2x_1} = b \oplus 0 \tag{6.21}$$

We may now try to construct φ^h_{xy} as $\bigwedge_{x'} \bigvee_{y: h(x) \geq h(y)} \psi^h_{x'y'}$. In the example, this gives $\varphi^h_{x_0 y_0} = (b \oplus 0) \wedge (c \oplus 0) \wedge (b \oplus 0) \equiv$ false, i.e., the interpretation is zero everywhere thereby not matching the map h. Applying the construction of Theorem 6.5.3 instead gives $\varphi^h_{x_0 y_0} = (b \oplus 0) \vee (c \oplus 0) \vee (b \oplus 0) \equiv (b \oplus 0) \vee (c \oplus 0)$, whose interpretation does match h. This also shows that we can not interchange the max and min in Lemma 6.5.2. ◂

Size of constructed formulas As discussed in the introduction, our constructions together yield an algorithm going from an approximation $\Gamma^i(\bot)(x, y)$ of the behavioural distance of states, via a proof tree, to a formula φ_{xy}. This is an alternative to the construction given as an algorithm in [RB23, Sec. 7]. In the worst case, this procedure will yield formulas whose size is exponential in the number of states of the corresponding LMC and the depth i of the approximation. We thus achieve the same asymptotic size complexity as the construction of *op. cit.*, however there are large classes of examples for which the optimisations in our proof system lead to smaller formulas (when counting the number of operators). The first example, taken from *op. cit.*, will show a notable improvement in size, and will allow us to discuss the shapes of LMCs leading to these improvements. Our final example applies to an LMC modelling random walks on the natural numbers. This is an infinite state example, which we are still able to capture as it is finitely branching.

Example 6.5.5: We will compare the size of formulas obtained via our construction with those obtained in an example of [RB23, Sec. 5]. The LMC involved can be represented as follows:

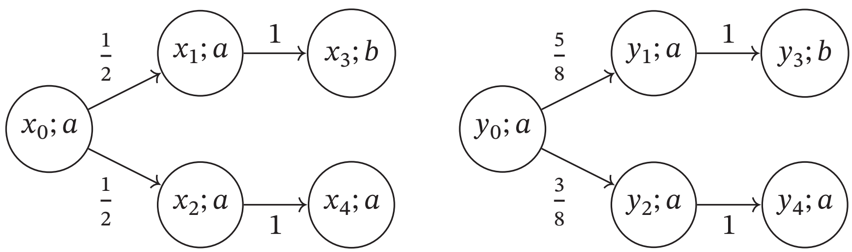

It can be computed that $\Gamma^3(\bot)(x_0, y_0) = \frac{1}{8}$. The construction applied in Theorem 6.3.9 gives the map $h_0\colon\ x_1, y_1 \mapsto 0, x_2, y_2 \mapsto 1$ for which $\tau(x_0) \cdot h_0 - \tau(y_0) \cdot h_0 = \frac{1}{8}$ so that we have:

$$\dfrac{\dfrac{\vdots}{x_2 \#_1 x_1} \quad \dfrac{\vdots}{x_2 \#_1 y_1} \quad \dfrac{\vdots}{y_2 \#_1 x_1} \quad \dfrac{\vdots}{y_2 \#_1 y_1}}{x_0 \#_{\frac{1}{8}} y_0}$$

The recursive lower bounds, given by $\Gamma^2(\bot)$, all have the same proof tree, up to renaming of states. For $x_2 \ \#_1 \ y_1$, we get $h_0 : x_4 \mapsto 1, y_3 \mapsto 0$ giving $\tau(x_4) \cdot h_0 - \tau(y_3) \cdot h_0 = 1$ so that:

$$\cfrac{\overline{x_4 \ \#_1 \ y_3}}{x_2 \ \#_1 \ y_1}$$

The formulas generated from such proofs are

$$\varphi_{x_2x_1} = \varphi_{x_2y_1} = \varphi_{y_2x_1} = \varphi_{y_2y_1} = \bigcirc\,[[(a \oplus 0) \wedge (\mathsf{false} \oplus 1)] \vee [(\mathsf{false} \oplus 0)]] \ominus 0 \tag{6.22}$$

Putting everything together, the formula $\varphi_{x_0y_0}$ is

$$\begin{aligned} \varphi_{x_0y_0} = \bigcirc[&[(\varphi_{x_2x_1} \oplus 0) \wedge (\varphi_{x_2y_1} \oplus 0) \wedge (\mathsf{false} \oplus 1)] \vee \\ &[(\varphi_{y_2x_1} \oplus 0) \wedge (\varphi_{y_2y_1} \oplus 0) \wedge (\mathsf{false} \oplus 1)] \vee \\ &[(\mathsf{false} \oplus 0)] \vee [(\mathsf{false} \oplus 0)]] \ominus \frac{3}{8} \end{aligned} \tag{6.23}$$

This has 8 recursive subformulas compared to the 100 occurring in the formula constructed in [RB23] and could all be written out within around 5 lines. Clearly, the formula is still not minimal; it can be simplified to $\bigcirc(\bigcirc a)) \ominus \frac{3}{8}$. This is because we can cancel conjunctions with false $\oplus$ 1 (which is everywhere 1) and disjunctions with false $\oplus$ 0 (which is everywhere 0). All recursive formulas are thus equivalent to $\bigcirc a$, and cancelling further yields the given formula. However, even without this simplification, we see a clear improvement in size. ◂

The main features which allow us to achieve such an improvement in the size of formulas witnessing lower bounds are: the number of states reachable at each step being less than the size of the entire state space; and those successors having non-zero behavioural distance so that the map h takes many different values. The combination of restricting to supports and omitting symmetric pairs from the recursive proof obligations in the *(exp)* rule, gives smaller proofs in these cases, which in turn are transformed into smaller formulas.

Example 6.5.6: We finish with an infinite state example based on random walks on the natural numbers. We model this as an LMC with state space $\mathbb{N}$ and transitions $\tau(n) = \frac{1}{2}\,|n-1\rangle + \frac{1}{2}\,|n+1\rangle$ for $n > 0$ and $\tau(0) = 1 \cdot |0\rangle$. Further, we have labels $\{a, b\}$ and labelling function $l(n) = \textbf{if}\ n = 0\ \textbf{then}\ b\ \textbf{else}\ a$.

States $n < m$ can clearly be distinguished by the probability to reach the state 0 with unique label b in n steps. In fact, this turns out to completely determine the distance between states. For example, $\Gamma^5(\bot)(4,6) = \frac{1}{2^4}$. This corresponds to the interpretation of the formula $\bigcirc^4 b$ on these states. In general, we have for $n < m$ and $i > n$, $\Gamma^i(\bot)(n,m) = \frac{1}{2^n}$. We will show the proof constructed for one such bound, as well as the formula constructed from this.

We have $\Gamma^3(\bot)(2,3) = \frac{1}{4}$ which can be proved as a lower bound using the map $h_0 \colon 1 \mapsto \frac{1}{2}, 3 \mapsto 0, 2 \mapsto 0, 4 \mapsto 0$ for which $\tau(2) \cdot h_0 - \tau(3) \cdot h_0 = \frac{1}{4}$ as follows

$$\dfrac{\dfrac{\vdots}{1 \,\#_{\frac{1}{2}}\, 2} \qquad \dfrac{\vdots}{1 \,\#_{\frac{1}{2}}\, 3} \qquad \dfrac{\vdots}{1 \,\#_{\frac{1}{2}}\, 4}}{2 \,\#_{\frac{1}{4}}\, 3}$$

The recursive proofs are all essentially the same; we show only the one for $1 \,\#_{\frac{1}{2}}\, 2$, which uses the map $h_0 \colon 0 \mapsto 1, 2 \mapsto 0, 1 \mapsto 0, 3 \mapsto 0$ yielding $\tau(1) \cdot h_0 - \tau(2) \cdot h_0 = \frac{1}{2}$ and

$$\dfrac{\dfrac{\vdots}{0 \,\#_1\, 1} \qquad \dfrac{\vdots}{0 \,\#_1\, 2} \qquad \dfrac{\vdots}{0 \,\#_1\, 3}}{1 \,\#_{\frac{1}{2}}\, 2}$$

For these recursive proofs, the corresponding formulas are

$$\varphi_{12} = \varphi_{13} = \varphi_{14} = \bigcirc[[(b \oplus 0) \wedge (b \oplus 0) \wedge (b \oplus 0) \wedge (\mathsf{false} \oplus 1)] \vee [\mathsf{false} \oplus 0] \vee [\mathsf{false} \oplus 0] \vee [\mathsf{false} \oplus 0]] \ominus 0 \tag{6.24}$$

From these, we obtain

$$\varphi_{23} = \bigcirc\Big[\Big[(\varphi_{12} \oplus 0) \wedge (\varphi_{13} \oplus 0) \wedge (\varphi_{14} \oplus 0) \wedge \Big(\mathsf{false} \oplus \frac{1}{2}\Big)\Big] \vee [\mathsf{false} \oplus 0] \vee [\mathsf{false} \oplus 0] \vee [\mathsf{false} \oplus 0]] \ominus 0 \tag{6.25}$$

This can be simplified to $\bigcirc[\bigcirc b \wedge (\mathsf{false} \oplus \frac{1}{2})]$, which is not equivalent to $\bigcirc\bigcirc b$, but does have the required property, namely that $|[\![\varphi_{23}]\!](2) - [\![\varphi_{23}]\!](3)| = \frac{1}{4}$. We thus obtain evidence for differences in the behaviour of states even in a system with an infinite state space. ◂

6.6 Conclusion

We have given a derivation system for lower bounds on behavioural distances between states in labelled Markov chains, with proofs of soundness and approximate completeness with respect to a least fixed point definition of the behavioural distance. The choice of the definition based on non-expansive maps was made specifically to allow the definition of a proof system, with the commonly used alternative definition based on couplings not immediately yielding a proof principle for lower bounds. The definitions are equivalent, by Kantorovich-Rubinstein duality [KR58]. This duality arises more generally when defining equivalences and their apartness counterparts via liftings, as was noted in the previous chapter. We further showed a close correspondence between proofs in our system and formulas in a modal logic, and compared this to the constructions in [RB23] going between finite approximations of distances and formulas in the same logic. We see quite some avenues for future work, and sketch some of them here.

Definitions of behavioural distances have been given for a variety of other system types, both metric and probabilistic, e.g., for Metric LTSs in [Bre05]; and Markov decision processes with finite [FPP04] and infinite state spaces [FPP05]. A natural extension would be to generalise our results in case we change the system type while keeping a similar definition of distance between distributions; however we may also consider adapted notions of distance, such as the total variation distance studied for LMCs in [CK14] or the MICo distance on MDPs [Cas+21]. Further statistical metrics/divergences such as the Lévy-Prokhorov metric [Pro56] or Kullback-Leibler divergence [KL51] may also be of interest. The former has recently [DS26] been shown to define a functional on pseudometric spaces, whose greatest fixed point is the ε-distance [DLT08].

A more general option would be to use the definition of codensity lifting [KSU18; Spr+21] and its suitability for capturing quantitative notions of equivalence to provide a sound and complete derivation system for many of these systems at once. The existing work on corresponding expressive logics [Kom+21] then gives us a starting point for providing a general version of the construction in Section 6.5. Another approach in the same vein is to use the theory of Kantorovich functors developed in [Gon+23], used to obtain characteristic logics also in the quantitative setting.

We would also like to investigate the connection to strategies in quantitative bisimulation games studied in [KM18; Vla+25] and developed for codensity bisimulations in [Kom+22].

In all these cases, the efficient computation of proofs and distinguishing formulas would be an interesting extension. There are a number of works in the qualitative setting (beyond the already discussed [RB23]) which could provide inspiration. For LTSs and branching bisimulation, computing (minimal) distinguishing formulas has been investigated by Martens and Groote [MG23; MG24]. König, Mika-Michalski and Schröder use coalgebraic techniques to develop algorithms for computing strategies in bisimulation games and transforming these into distinguishing formulas [KMS20]. Wißmann, Milius and Schröder give a coalgebraic algorithm related to partition refinement which constructs modal formulas characterising behavioural equivalence classes [WMS22].

An alternative approach to improving the robustness of probabilistic bisimilarity, are approximate bisimulations, such as: ε-bisimilarity [DLT08]; ε-APB [DAK12]; and ε-lumpability [Buc94]. These are often close to existing qualitative definitions, with some degree of error introduced. For a recent overview, and extension to weak and branching bisimulation, see [Spo+24]. It would be interesting to further compare these approximate notions to distances, as has been done for ε-bisimilarity in [DS26], and relate them to proofs and logics.

Part III

Coda

CHAPTER 7

Conclusions

IN THE FIRST PART OF THIS THESIS, we have seen how the effect on behaviours of constructions on coalgebras can be analysed in terms of their transportation of coinductive predicates. Our starting point was two coalgebraic generalisations of automata constructions: determinisation of probabilistic automata; and ultrafilter extensions. These are both based on composing liftings of adjoint functors to categories of coalgebras, first adding algebraic/topological structure, and then forgetting it.

Such liftings are outside the scope of earlier work, as they are not induced by the distributive laws assumed there. Instead, in Chapter 3, we considered liftings induced by what we called *invertible steps*. In the examples of determinisation and ultrafilter extensions, these arose from *weak* distributive laws and the corresponding *weak* liftings.

The other main ingredient of our analysis, is the lifting of the adjoint functors underlying the coalgebraic constructions to fibrations. By relating the liftings to coalgebras and fibrations via inequalities, we obtained conditions for the preservation and reflection of (co)inductive invariants and predicates. The applicability of these conditions was shown by:

- proving that forgetting algebraic structure on coalgebras, as is done in the second stage of coalgebraic constructions of determinised automata and ultrafilter extensions, preserves and reflects bisimilarity/coalgebraic behavioural equivalence;
- recovering results relating bisimulations on Kripke frames and Vietoris frames, including that they coincide after forgetting topological structure;

In Chapter 4, we generalised our approach further, by considering the liftings of left and right adjoints separately. The *invertible* steps we assumed in the previous chapter, required the *forward* and *backward* steps to be each other's one-sided inverses. Steps coming from weak liftings satisfy this requirement, but one account of the determinisation of probabilistic automata already gives an example where it can not be guaranteed. That approach instead assumes a *lax* lifting, induced by a step which may not have a *natural* inverse. Our generalised conditions allowed us to show preservation and reflection of bisimilarity also in this setting.

The generality of these new conditions also allowed applications to the dual adjunction-based setting of coalgebraic modal logic. There, the semantics of modal formulas are given by only a *forward* step. Our conditions involving only these types of step instantiate to relations between bisimilarity and logical equivalence capturing adequacy and expressivity of the logic. Alternatively, our conditions can be used to relate apartness and the existence of distinguishing formulas to give a dual view of adequacy and expressivity.

In the second part of this thesis, we further investigated apartness as a notion of distinguishability of behaviours. To facilitate the inductive derivation of differences in behaviour of systems involving probabilistic transitions, we dualised the definition of precongruences due to Aczel and Mendler [AM89]. This yielded an opposite notion to coalgebraic behavioural equivalence, which, in contrast to an earlier definition, could be proved in probabilistic systems by providing finite evidence. To avoid a restriction to systems with finite state spaces, we further showed that attention can be restricted to one-step-reachable states. The correspondence with our definition of behavioural apartness was guaranteed by proofs of soundness and completeness.

Continuing the development of means for distinguishing behaviours of probabilistic systems, we then moved to a form of apartness arising from behavioural distances. We showed that an existing definition of behavioural distance involving a non-expansive-map-based lifting, allows for finite proofs of lower bounds on the distance between states. The correspondence with the definition of behavioural distance was more subtle than the case of two-valued behavioural apartness. Soundness is as expected, but we can not show "true" completeness: there are systems in which the distance of states can not be proved exactly in a finite number of steps. Instead, we showed approximate completeness: that we can prove bounds arbitrarily close to the true value. Finally, we showed a constructive correspondence between proofs and formulas of a rational-valued modal logic, used in earlier work for "explanations" of behavioural distances [RB23], leading to smaller formulas in classes of examples.

Connections and Future Work

Though we have presented this thesis in two parts, each aims to understand the behaviour of systems which can be modelled by coalgebras. At a more technical level, the connections between the parts and chapters have become clearer during the work of this thesis. We sketch the relations here, allowing us to formulate some open questions we hope will be revisited in the future.

As already discussed with respect to our definition of behavioural apartness, and alluded to for the derivation system for lower bounds, both can be viewed via liftings to fibrations. The comparison of apartness and existence of distinguishing formulas, as we did for a specific modal logic in Chapter 4, should then be generalisable to these new notions. One interesting question is then if we can say something constructive about these comparisons. For instance, can we extract procedures for building distinguishing formulas directly from the fibrational theory? A missing ingredient for this seems to be a fibrational definition of apartness *proofs*, which should be related to the existing definitions of modal formulas and their semantics.

An alternative approach to a more general view of distinguishability in coalgebras may be given by *codensity lifting* [KSU18]. Its general suitability for quantitative applications [Spr+21; Kom+21] warrant further study related to the results of this thesis.

On the one hand, it would be interesting to apply our conditions of Chapters 3 and 4 to codensity liftings. This requires a definition of codensity lifting of functors between different categories, analogous to the predicate, relation, and quotient liftings used in Chapter 4. Such a generalisation has recently been given in [Kor+24]. Also related is the "transfer of codensity bisimilarity" result provided by Komorida et al. in [Kom+22]. This relates codensity bisimilarity for codensity liftings of a single endofunctor to two different fibrations via a functor between the fibrations, already in a similar spirit to our results in Chapter 4.

On the other hand, it has been shown that codensity lifting generalises the non-expansive map based lifting we used in Chapter 6, and should allow witnesses of a similar form in apartness proofs (as we discuss further below). We would like to have a complete picture of the relationships of bisimilarity and apartness, with modal logic and games, within the scope of codensity lifting, described next.

Codensity Apartness

Some of the connections between equivalences defined using codensity lifting, modal logics, and games have already been shown in the literature: a development of expressive logics arising from the modalities present in codensity lifting has been given in [Kom+21], and a game characterisation of codensity bisimilarity has been given in [Kom+22].

Outside the scope of codensity lifting, the constructive picture for labelled transition systems is quite complete. The construction of distinguishing formulas for apart states was given by Geuvers [Geu22]. The relation between apartness and winning strategies for the spoiler in bisimulation games has been shown in [RJB24; KW24], although non-constructive arguments are used to cover infinite state systems.

We would like to have a similar account for (in)equivalences defined via codensity lifting to have correspondences between: proofs of apartness; distinguishing formulas; and spoiler strategies in games. While the definitions of logics and games have already been given, a good definition of apartness is a missing ingredient. The two directions taken in Chapters 5 and 6 give possible starting points.

The first suggests the definition of an opposite codensity lifting, yielding a dual definition of codensity *co*bisimilarity. On the other hand, one could work directly with the existing definition of the lifting as we did for the behavioural distance on LMCs, and use a similar proof technique.

The usability of each approach seems to depend on the objects of interest. Some clarity can be gained from considering the proof technique we have used to obtain lower bounds in Chapter 6 more abstractly. In each step of a proof, our inductive hypothesis is a distance which is below the true behavioural distance. Then, we apply the lifting in the definition of behavioural distance to this hypothesis, and a lower bound on this gives a lower bound on the true distance. We can write this one-step rule abstractly as:

$$\frac{d \leq \mathrm{lfp}(\Gamma) \qquad d' \leq \Gamma(d)}{d' \leq \mathrm{lfp}(\Gamma)} \tag{7.1}$$

Soundness follows by observing the following chain of inequalities: $d' \leq \Gamma(d) \leq \Gamma(\mathrm{lfp}(\Gamma)) = \mathrm{lfp}(\Gamma)$.

We can also dualise the above proof principle, to give a means of establishing elements to be above the greatest fixed point:

$$\frac{\mathrm{gfp}(\varphi) \leq p \qquad \varphi(p) \leq p'}{\mathrm{gfp}(\varphi) \leq p'} \tag{7.2}$$

Each principle can be chosen to match the type of fixpoint being reasoned about. In either case, the liftings we have used are well-suited to this reasoning, as they allow us to quite straightforwardly establish inequalities $p' \leq \varphi(p)$ or $\varphi(p) \leq p'$. As may be expected, codensity lifting also fits this view. It is defined by a meet ranging over a collection of maps, sometimes called "tests". A choice of test thus shows an inequality of the form $\varphi(p) \leq p'$, usable in the rule for greatest fixed points.

It is known that an instance of codensity lifting recovers the behavioural distance of Chapter 6 [Kom+21]. This is a situation where the opposite lifting approach may be more intuitive, as that lifting avoids the need to reverse the ordering most often taken in the literature.

We are left with the following questions:

1. To which settings can we apply opposite codensity lifting?
2. Can we connect proofs of apartness notions derived from codensity liftings to distinguishing formulas in the corresponding modal logics?
3. Do these correspondences also extend to (spoiler) strategies in codensity games?
4. Can we obtain (general) procedures for moving between these different forms of evidence?

Algorithms

The algorithms given in Chapter 6 could be implemented to automatically provide our construction of distinguishing formulas for labelled Markov chains. Increasing the efficiency of such implementations may also lead to improvements of the apartness proof systems. The main area of improvement would seem to be the computation of optimal maps used as witnesses in proofs. We have already discussed the possibility of restricting the search space to vertices of polytopes of non-expansive maps. Perhaps further restrictions are possible, justified by properties of convex optimisation.

Related to this, is a more general question of proof search in our system. Our completeness proof suggests an algorithm for constructing proofs of the n-step distance between states for any n. The $N+1$-th step of such a proof (using the *(exp)* rule) relies on recursive proofs of the exact N-step distances of successor states in order to construct a map h. We would like to investigate the possibility of finding good candidates for the map h given just a desired lower bound (and of course the LMC of interest).

For our purposes in Chapter 6, the chosen rules were a good fit to keep the complexity of proofs and constructions done by hand under control. However, in the context of automation, it may be useful to incorporate further rules, for instance more rules dual to those developed in work on quantitative equational theories [MPP16]. These additions would make the choice of rule at each stage of a proof a larger focus of proof search. Existing work on so-called goal directed proof search, which has proven effective for, e.g., linear logic [HM91; CHP00] and separation logics [Chl11; Sam+21], could then be of interest. This approach aims to generate proof steps bottom-up looking solely at the shape of the proof goal.

The more recently introduced *incorrectness* (separation) logics [OHe20; Raa+20] may also serve as inspiration, as they deal with proofs of the existence of bugs in programs rather than their absence, in a similar spirit to apartness. This has been used within the tool Infer [Dis+19] to detect and report memory safety bugs in C++ codebases.

Further logical properties and constructions

Our results on the construction of ultrafilter extensions for coalgebras did not concern their "bisimilarity-somewhere-else" property, for which they were developed in the study of modal logic [BRV01]. Namely, we did not show that after the first step of the construction, where the ultrafilter structure is added, logically equivalent states are also bisimilar/behaviourally equivalent. In the original work providing an ultrafilter constructions for coalgebras [KKP05], the behavioural equivalence of logically equivalent states is established. Considering our later application to expressivity of logics, it would be interesting to see if our conditions involving steps and liftings can be used to say more about the logical properties of ultrafilter extensions.

Closely related to expressivity and the correspondence between apartness and distinguishing formulas, is the concept of characteristic formulas. These are single formulas which are satisfied by exactly those states in an equivalence class of some notion of behavioural equivalence. It seems possible that the existence of such formulas could be determined by conditions related to those of expressivity. Existing works on constructing distinguishing formulas based on partition refinement in fact already provide such formulas [WMS22]. We would like to understand how this may relate to our apartness proofs and related constructions.

Bibliography

[AM89] Peter Aczel and Nax Paul Mendler. 'A Final Coalgebra Theorem'. In: *Category Theory and Computer Science*. Vol. 389. Lecture Notes in Computer Science. Springer, 1989, pp. 357–365 (cit. on pp. 10, 61, 98, 104, 150).

[AHS09] Jirí Adámek, Horst Herrlich and George E. Strecker. *Abstract and Concrete Categories - The Joy of Cats*. Dover Publications, 2009. ISBN: 978-0-486-46934-8 (cit. on p. 27).

[Ash72] Robert B Ash. *Real Analysis and Probability: Probability and Mathematical Statistics: A Series of Monographs and Textbooks*. Academic press, 1972 (cit. on p. 137).

[Bac+17] Giorgio Bacci, Giovanni Bacci, Kim G. Larsen and Radu Mardare. 'On-the-Fly Computation of Bisimilarity Distances'. In: *Log. Methods Comput. Sci.* 13.2 (2017) (cit. on p. 120).

[Bac+18] Giorgio Bacci, Giovanni Bacci, Kim G. Larsen and Radu Mardare. 'A Complete Quantitative Deduction System for the Bisimilarity Distance on Markov Chains'. In: *Log. Methods Comput. Sci.* 14.4 (2018) (cit. on p. 123).

[Bac+21] Giorgio Bacci, Giovanni Bacci, Kim G. Larsen, Radu Mardare, Qiyi Tang and Franck van Breugel. 'Computing Probabilistic Bisimilarity Distances for Probabilistic Automata'. In: *Log. Methods Comput. Sci.* 17.1 (2021) (cit. on p. 120).

[Bal+18] Paolo Baldan, Filippo Bonchi, Henning Kerstan and Barbara König. 'Coalgebraic Behavioral Metrics'. In: *Log. Methods Comput. Sci.* 14.3 (2018) (cit. on pp. 6, 120, 122).

[Bal+23] Paolo Baldan, Richard Eggert, Barbara König and Tommaso Padoan. 'Fixpoint Theory - Upside Down'. In: *Log. Methods Comput. Sci.* 19.2 (2023) (cit. on pp. 120, 122).

[BKR19] Simone Barlocco, Clemens Kupke and Jurriaan Rot. 'Coalgebra Learning via Duality'. In: *FoSSaCS*. Vol. 11425. Lecture Notes in Computer Science. Springer, 2019, pp. 62–79 (cit. on p. 110).

[Bar71] Michael Barr. 'Exact categories'. In: *Exact Categories and Categories of Sheaves*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1971, pp. 1–120. ISBN: 978-3-540-36999-8 (cit. on p. 58).

[BSV04] Falk Bartels, Ana Sokolova and de Vink,Erik P. 'A hierarchy of probabilistic system types'. In: *Theor. Comput. Sci.* 327.1-2 (2004), pp. 3–22 (cit. on pp. 55, 105).

[Bas18] Henning Basold. 'Mixed Inductive-Coinductive Reasoning Types, Programs and Logic'. PhD thesis. Radboud Universiteit Nijmegen, 2018 (cit. on pp. 13, 98).

[Bec06] Jon Beck. 'Distributive laws'. In: *Seminar on Triples and Categorical Homology Theory: ETH 1966/67*. Springer. 2006, pp. 119–140 (cit. on p. 21).

[Ben79] Johan van Benthem. 'Canonical Modal Logics and Ultrafilter Extensions'. In: *J. Symb. Log.* 44.1 (1979), pp. 1–8 (cit. on p. 43).

[Beo+23] Harsh Beohar, Sebastian Gurke, Barbara König and Karla Messing. 'Hennessy-Milner Theorems via Galois Connections'. In: *CSL*. Vol. 252. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, 12:1–12:18 (cit. on pp. 12, 67, 69, 93, 117).

[Beo+24] Harsh Beohar, Sebastian Gurke, Barbara König, Karla Messing, Jonas Forster, Lutz Schröder and Paul Wild. 'Expressive Quantale-Valued Logics for Coalgebras: An Adjunction-Based Approach'. In: *STACS*. Vol. 289. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024, 10:1–10:19 (cit. on pp. 12, 117).

[BFV10] Nick Bezhanishvili, Gaëlle Fontaine and Yde Venema. 'Vietoris Bisimulations'. In: *J. Log. Comput.* 20.5 (2010), pp. 1017–1040 (cit. on pp. 35, 36, 45, 58–60, 62, 63, 68, 69).

[BRV01] Patrick Blackburn, de Rijke,Maarten and Yde Venema. *Modal Logic*. Vol. 53. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001. ISBN: 978-1-10705088-4 (cit. on pp. 35, 43, 154).

[Blo12] Alwin Blok. 'Interaction, observation and denotation'. MSc Thesis. Universiteit van Amsterdam, 2012 (cit. on p. 110).

[Böh10] Gabriella Böhm. 'The weak theory of monads'. In: *Advances in Mathematics* 225.1 (2010), pp. 1–32 (cit. on p. 21).

[Bon+12] Filippo Bonchi, Marcello M. Bonsangue, Michele Boreale, Jan J. M. M. Rutten and Alexandra Silva. 'A coalgebraic perspective on linear weighted automata'. In: *Inf. Comput.* 211 (2012), pp. 77–105 (cit. on p. 33).

[Bon+16] Filippo Bonchi, Marcello M. Bonsangue, Georgiana Caltais, Jan Rutten and Alexandra Silva. 'A coalgebraic view on decorated traces'. In: *Math. Struct. Comput. Sci.* 26.7 (2016), pp. 1234–1268 (cit. on p. 33).

[BKP23] Filippo Bonchi, Barbara König and Daniela Petrisan. 'Up-to techniques for behavioural metrics via fibrations'. In: *Math. Struct. Comput. Sci.* 33.4-5 (2023), pp. 182–221 (cit. on pp. 6, 29, 120, 122).

[Bon+17] Filippo Bonchi, Daniela Petrisan, Damien Pous and Jurriaan Rot. 'A general account of coinduction up-to'. In: *Acta Informatica* 54.2 (2017), pp. 127–190 (cit. on pp. 33, 50, 122).

[BS21] Filippo Bonchi and Alessio Santamaria. 'Combining Semilattices and Semimodules'. In: *FoSSaCS*. Vol. 12650. Lecture Notes in Computer Science. Springer, 2021, pp. 102–123 (cit. on pp. 34, 76).

[BSS17] Filippo Bonchi, Alexandra Silva and Ana Sokolova. 'The Power of Convex Algebras'. In: *CONCUR*. Vol. 85. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017, 23:1–23:18 (cit. on pp. 38, 55).

[BSS21] Filippo Bonchi, Alexandra Silva and Ana Sokolova. 'Distribution Bisimilarity via the Power of Convex Algebras'. In: *Log. Methods Comput. Sci.* 17.3 (2021) (cit. on pp. 7, 34–36, 38, 41, 45, 55, 60, 61, 66, 67, 69, 76, 83).

[BK05] Marcello M. Bonsangue and Alexander Kurz. 'Duality for Logics of Transition Systems'. In: *FoSSaCS*. Vol. 3441. Lecture Notes in Computer Science. Springer, 2005, pp. 455–469 (cit. on pp. 36, 65).

[BMS13] Marcello M. Bonsangue, Stefan Milius and Alexandra Silva. 'Sound and Complete Axiomatizations of Coalgebraic Language Equivalence'. In: *ACM Trans. Comput. Log.* 14.1 (2013), 7:1–7:52 (cit. on pp. 33, 65).

[Bor94a] Francis Borceux. *Handbook of categorical algebra: volume 1, Basic category theory*. Vol. 1. Cambridge University Press, 1994 (cit. on p. 28).

[Bor94b] Francis Borceux. *Handbook of categorical algebra: volume 2, Categories and Structures.* Vol. 2. Cambridge University Press, 1994 (cit. on pp. 20, 28, 51, 58).

[Bre05] Franck van Breugel. 'A Behavioural Pseudometric for Metric Labelled Transition Systems'. In: *CONCUR*. Vol. 3653. Lecture Notes in Computer Science. Springer, 2005, pp. 141–155 (cit. on p. 144).

[Bre12] Franck van Breugel. 'On behavioural pseudometrics and closure ordinals'. In: *Inf. Process. Lett.* 112.19 (2012), pp. 715–718 (cit. on p. 125).

[Bre17] Franck van Breugel. 'Probabilistic bisimilarity distances'. In: *ACM SIGLOG News* 4.4 (2017), pp. 33–51 (cit. on pp. 119, 124).

[Bre+07] Franck van Breugel, Claudio Hermida, Michael Makkai and James Worrell. 'Recursively defined metric spaces without contraction'. In: *Theor. Comput. Sci.* 380.1-2 (2007), pp. 143–163 (cit. on p. 121).

[BSW08] Franck van Breugel, Babita Sharma and James Worrell. 'Approximating a Behavioural Pseudometric without Discount for Probabilistic Systems'. In: *Log. Methods Comput. Sci.* 4.2 (2008) (cit. on p. 119).

[BW05] Franck van Breugel and James Worrell. 'A behavioural pseudometric for probabilistic transition systems'. In: *Theor. Comput. Sci.* 331.1 (2005), pp. 115–142 (cit. on p. 119).

[Bro05] James Brotherston. 'Cyclic Proofs for First-Order Logic with Inductive Definitions'. In: *TABLEAUX*. Vol. 3702. Lecture Notes in Computer Science. Springer, 2005, pp. 78–92 (cit. on p. 12).

[Buc94] Peter Buchholz. 'Exact and ordinary lumpability in finite Markov chains'. In: *Journal of applied probability* 31.1 (1994), pp. 59–75 (cit. on p. 145).

[Cas+21] Pablo Samuel Castro, Tyler Kastner, Prakash Panangaden and Mark Rowland. 'MICo: Improved representations via sampling-based state similarity for Markov decision processes'. In: *NeurIPS*. 2021, pp. 30113–30126 (cit. on p. 144).

[CC95] Ufuk Celikkan and Rance Cleaveland. 'Generating Diagnostic Information for Behavioral Preorders'. In: *Distributed Comput.* 9.2 (1995), pp. 61–75 (cit. on p. 13).

[CHP00] Iliano Cervesato, Joshua S. Hodas and Frank Pfenning. 'Efficient resource management for linear logic proof search'. In: *Theor. Comput. Sci.* 232.1-2 (2000), pp. 133–163 (cit. on p. 154).

[CBW12] Di Chen, Franck van Breugel and James Worrell. 'On the Complexity of Computing Probabilistic Bisimilarity'. In: *FoSSaCS*. Vol. 7213. Lecture Notes in Computer Science. Springer, 2012, pp. 437–451 (cit. on p. 119).

[CJ14] Liang-Ting Chen and Achim Jung. 'On a Categorical Framework for Coalgebraic Modal Logic'. In: *MFPS*. Vol. 308. Electronic Notes in Theoretical Computer Science. Elsevier, 2014, pp. 109–128 (cit. on pp. 36, 65).

[CK14] Taolue Chen and Stefan Kiefer. 'On the total variation distance of labelled Markov chains'. In: *CSL-LICS*. ACM, 2014, 33:1–33:10 (cit. on p. 144).

[Chl11] Adam Chlipala. 'Mostly-automated verification of low-level programs in computational separation logic'. In: *PLDI*. ACM, 2011, pp. 234–245 (cit. on p. 154).

[Cle90] Rance Cleaveland. 'On Automatically Explaining Bisimulation Inequivalence'. In: *CAV*. Vol. 531. Lecture Notes in Computer Science. Springer, 1990, pp. 364–372 (cit. on p. 13).

[Clo+16] Ranald Clouston, Ales Bizjak, Hans Bugge Grathwohl and Lars Birkedal. 'The Guarded Lambda-Calculus: Programming and Reasoning with Guarded Recursion for Coinductive Types'. In: *Log. Methods Comput. Sci.* 12.3 (2016) (cit. on p. 98).

[Coq93] Thierry Coquand. 'Infinite Objects in Type Theory'. In: *TYPES*. Vol. 806. Lecture Notes in Computer Science. Springer, 1993, pp. 62–78 (cit. on p. 12).

[DAK12] Alessandro D'Innocenzo, Alessandro Abate and Joost-Pieter Katoen. 'Robust PCTL model checking'. In: *HSCC*. ACM, 2012, pp. 275–286 (cit. on p. 145).

[DP02] Brian A. Davey and Hilary A. Priestley. *Introduction to Lattices and Order, Second Edition*. Cambridge University Press, 2002 (cit. on p. 46).

[DEP98] Josée Desharnais, Abbas Edalat and Prakash Panangaden. 'A Logical Characterization of Bisimulation for Labeled Markov Processes'. In: *LICS*. IEEE Computer Society, 1998, pp. 478–487 (cit. on pp. 11, 121).

[DEP02] Josée Desharnais, Abbas Edalat and Prakash Panangaden. 'Bisimulation for Labelled Markov Processes'. In: *Inf. Comput.* 179.2 (2002), pp. 163–193 (cit. on pp. 11, 121).

[Des+99] Josée Desharnais, Vineet Gupta, Radha Jagadeesan and Prakash Panangaden. 'Metrics for Labeled Markov Systems'. In: *CONCUR*. Vol. 1664. Lecture Notes in Computer Science. Springer, 1999, pp. 258–273 (cit. on pp. 11, 119, 121).

[Des+04] Josée Desharnais, Vineet Gupta, Radha Jagadeesan and Prakash Panangaden. 'Metrics for labelled Markov processes'. In: *Theor. Comput. Sci.* 318.3 (2004), pp. 323–354 (cit. on p. 134).

[Des+02] Josée Desharnais, Radha Jagadeesan, Vineet Gupta and Prakash Panangaden. 'The Metric Analogue of Weak Bisimulation for Probabilistic Processes'. In: *LICS*. IEEE Computer Society, 2002, pp. 413–422 (cit. on p. 125).

[DLT08] Josée Desharnais, François Laviolette and Mathieu Tracol. 'Approximate Analysis of Probabilistic Processes: Logic, Simulation and Games'. In: *QEST*. IEEE Computer Society, 2008, pp. 264–273 (cit. on pp. 144, 145).

[DS26] Josée Desharnais and Ana Sokolova. 'ϵ-Distance via Lévy-Prokhorov Lifting'. In: *CSL*. Vol. 363. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2026, 26:1–26:24 (cit. on pp. 144, 145).

[Dis+19] Dino Distefano, Manuel Fähndrich, Francesco Logozzo and Peter W. O'Hearn. 'Scaling static analyses at Facebook'. In: *Commun. ACM* 62.8 (2019), pp. 62–70 (cit. on p. 154).

[Dob06] Ernst-Erich Doberkat. 'Eilenberg-Moore algebras for stochastic relations'. In: *Inf. Comput.* 204.12 (2006), pp. 1756–1781 (cit. on p. 20).

[ES18] Sebastian Enqvist and Sumit Sourabh. 'Bisimulations for coalgebras on Stone spaces'. In: *J. Log. Comput.* 28.6 (2018), pp. 991–1010 (cit. on p. 63).

[FPP04] Norm Ferns, Prakash Panangaden and Doina Precup. 'Metrics for Finite Markov Decision Processes'. In: *AAAI*. AAAI Press/The MIT Press, 2004, pp. 950–951 (cit. on p. 144).

[FPP05] Norm Ferns, Prakash Panangaden and Doina Precup. 'Metrics for Markov Decision Processes with Infinite State Spaces'. In: *UAI*. AUAI Press, 2005, pp. 201–208 (cit. on p. 144).

[FKP17] Nathanaël Fijalkow, Bartek Klin and Prakash Panangaden. 'Expressiveness of Probabilistic Modal Logics, Revisited'. In: *ICALP*. Vol. 80. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017, 105:1–105:12 (cit. on pp. 11, 117, 121).

[For+23] Jonas Forster, Sergey Goncharov, Dirk Hofmann, Pedro Nora, Lutz Schröder and Paul Wild. 'Quantitative Hennessy-Milner Theorems via Notions of Density'. In: *CSL*. Vol. 252. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, 22:1–22:20 (cit. on pp. 12, 94, 117).

[Gar20] Richard Garner. 'The Vietoris Monad and Weak Distributive Laws'. In: *Appl. Categorical Struct.* 28.2 (2020), pp. 339–354 (cit. on pp. 8, 21, 22, 34, 35, 37, 38, 40, 76).

[Geu22] Herman Geuvers. 'Apartness and Distinguishing Formulas in Hennessy-Milner Logic'. In: *A Journey from Process Algebra via Timed Automata to Model Learning*. Vol. 13560. Lecture Notes in Computer Science. Springer, 2022, pp. 266–282 (cit. on pp. 67, 90, 91, 117, 122, 152).

[GJ21] Herman Geuvers and Bart Jacobs. 'Relating Apartness and Bisimulation'. In: *Log. Methods Comput. Sci.* 17.3 (2021) (cit. on pp. 9, 10, 67, 90, 91, 97–102, 114, 120, 122).

[GJS90] Alessandro Giacalone, Chi-Chang Jou and Scott A. Smolka. 'Algebraic Reasoning for Probabilistic Concurrent Systems'. In: *Programming Concepts and Methods*. North-Holland, 1990, pp. 443–458 (cit. on pp. 10, 119).

[GL95] Jens Chr. Godskesen and Kim Guldstrand Larsen. 'Synthesizing Distinguishing Formulae for Real Time Systems'. In: *Nord. J. Comput.* 2.3 (1995), pp. 338–357 (cit. on p. 13).

[GLR00] Joseph A. Goguen, Kai Lin and Grigore Rosu. 'Circular Coinductive Rewriting'. In: *ASE*. IEEE Computer Society, 2000, pp. 123–132 (cit. on p. 12).

[Gol74] Robert Ian Goldblatt. 'Metamathematics of modal logic'. In: *Bulletin of the Australian Mathematical Society* 10.3 (1974). Publisher: Cambridge University Press, pp. 479–480 (cit. on p. 43).

[Gon+23] Sergey Goncharov, Dirk Hofmann, Pedro Nora, Lutz Schröder and Paul Wild. 'Kantorovich Functors and Characteristic Logics for Behavioural Distances'. In: *FoSSaCS*. Vol. 13992. Lecture Notes in Computer Science. Springer, 2023, pp. 46–67 (cit. on pp. 12, 94, 117, 144).

[Goy21] Alexandre Goy. 'On the compositionality of monads via weak distributive laws. (Compositionnalité des monades par lois de distributivité faibles)'. PhD thesis. University of Paris-Saclay, France, 2021 (cit. on pp. 7, 22, 34, 36, 38–41, 44, 45, 76).

[GP20] Alexandre Goy and Daniela Petrisan. 'Combining probabilistic and non-deterministic choice via weak distributive laws'. In: *LICS*. ACM, 2020, pp. 454–464 (cit. on pp. 7, 34, 36, 38, 41, 44, 45, 76).

[GPA21] Alexandre Goy, Daniela Petrisan and Marc Aiguier. 'Powerset-Like Monads Weakly Distribute over Themselves in Toposes and Compact Hausdorff Spaces'. In: *ICALP*. Vol. 198. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 132:1–132:14 (cit. on p. 34).

[Gra21] Marino Gran. 'An introduction to regular categories'. In: *New Perspectives in Algebra, Topology and Categories: Summer School, Louvain-la-Neuve, Belgium, September 12-15, 2018 and September 11-14, 2019* (2021), pp. 113–145 (cit. on p. 23).

[Gro22] Jim de Groot. 'Dualities in modal logic'. PhD thesis. 2022 (cit. on p. 66).

[Gum99] H. Peter Gumm. 'Elements of the general theory of coalgebras'. In: *LUATCS 99*. Rand Afrikaans University, South Africa, 1999 (cit. on p. 104).

[Gum09] H. Peter Gumm. 'Copower functors'. In: *Theor. Comput. Sci.* 410.12-13 (2009), pp. 1129–1142 (cit. on p. 115).

[GS01] H. Peter Gumm and Tobias Schröder. 'Monoid-labeled transition systems'. In: *CMCS*. Vol. 44. Electronic Notes in Theoretical Computer Science 1. Elsevier, 2001, pp. 185–204 (cit. on p. 115).

[GT22] H. Peter Gumm and Mona Taheri. 'Saturated Kripke Structures as Vietoris Coalgebras'. In: *CMCS*. Vol. 13225. Lecture Notes in Computer Science. Springer, 2022, pp. 88–109 (cit. on p. 63).

[HJS07] Ichiro Hasuo, Bart Jacobs and Ana Sokolova. 'Generic Trace Semantics via Coinduction'. In: *Log. Methods Comput. Sci.* 3.4 (2007) (cit. on p. 65).

[HKC18] Ichiro Hasuo, Toshiki Kataoka and Kenta Cho. 'Coinductive predicates and final sequences in a fibration'. In: *Math. Struct. Comput. Sci.* 28.4 (2018), pp. 562–611 (cit. on pp. 28, 29, 63, 66, 108).

[HMS05] Daniel Hausmann, Till Mossakowski and Lutz Schröder. 'Iterative Circular Coinduction for CoCasl in Isabelle/HOL'. In: *FASE*. Vol. 3442. Lecture Notes in Computer Science. Springer, 2005, pp. 341–356 (cit. on p. 12).

[Her92] Claudio Hermida. 'On Fibred Adjunctions and Completeness for Fibred Categories'. In: *COMPASS/ADT*. Vol. 785. Lecture Notes in Computer Science. Springer, 1992, pp. 235–251 (cit. on pp. 50, 70, 72, 74).

[Her93] Claudio Hermida. 'Fibrations, logical predicates and indeterminates'. PhD thesis. University of Edinburgh, UK, 1993 (cit. on pp. 5, 66, 70–72, 74, 75).

[HJ98] Claudio Hermida and Bart Jacobs. 'Structural Induction and Coinduction in a Fibrational Setting'. In: *Inf. Comput.* 145.2 (1998), pp. 107–152 (cit. on pp. 5, 37, 42, 47, 66, 99, 120).

[Hey66] Arend Heyting. *Intuitionism: an introduction*. Vol. 41. Elsevier, 1966 (cit. on p. 120).

[HM91] Joshua S. Hodas and Dale Miller. 'Logic Programming in a Fragment of Intuitionistic Linear Logic'. In: *LICS*. IEEE Computer Society, 1991, pp. 32–42 (cit. on p. 154).

[HMU07] John E. Hopcroft, Rajeev Motwani and Jeffrey D. Ullman. *Introduction to automata theory, languages, and computation, 3rd Edition*. Pearson international edition. Addison-Wesley, 2007 (cit. on p. 2).

[HJ04] Jesse Hughes and Bart Jacobs. 'Simulations in coalgebra'. In: *Theor. Comput. Sci.* 327.1-2 (2004), pp. 71–108 (cit. on p. 29).

[Jac16] Bart Jacobs. *Introduction to Coalgebra: Towards Mathematics of States and Observation*. Vol. 59. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2016 (cit. on pp. 4, 23, 26, 28, 50, 52, 62, 114).

[Jac] Bart Jacobs. 'Structured Probabilistic Reasoning'. Unpublished draft, accessed 2024-10-30. Unpublished draft (cit. on p. 123).

[JSS15] Bart Jacobs, Alexandra Silva and Ana Sokolova. 'Trace semantics via determinization'. In: *J. Comput. Syst. Sci.* 81.5 (2015), pp. 859–879 (cit. on pp. 7, 33, 37, 44, 65).

[JS10] Bart Jacobs and Ana Sokolova. 'Exemplaric Expressivity of Modal Logics'. In: *J. Log. Comput.* 20.5 (2010), pp. 1041–1068 (cit. on pp. 69, 76).

[Jac01] Bart P. F. Jacobs. *Categorical Logic and Type Theory*. Vol. 141. Studies in logic and the foundations of mathematics. North-Holland, 2001 (cit. on pp. 25, 50, 52, 66, 70, 72, 75, 91).

[KS90] Paris C. Kanellakis and Scott A. Smolka. 'CCS Expressions, Finite State Processes, and Three Problems of Equivalence'. In: *Inf. Comput.* 86.1 (1990), pp. 43–68 (cit. on p. 13).

[KR58] Leonid Vasilevich Kantorovich and SG Rubinshtein. 'On a space of totally additive functions'. In: *Vestnik of the St. Petersburg University: Mathematics* 13.7 (1958), pp. 52–59 (cit. on p. 144).

[KSU18] Shin-ya Katsumata, Tetsuya Sato and Tarmo Uustalu. 'Codensity Lifting of Monads and its Dual'. In: *Log. Methods Comput. Sci.* 14.4 (2018) (cit. on pp. 117, 144, 151).

[KW24] Jeroen J. A. Keiren and Tim A. C. Willemse. 'It's All a Game - Apartness and Bisimilarity'. In: *Logics and Type Systems in Theory and Practice*. Vol. 14560. Lecture Notes in Computer Science. Springer, 2024, pp. 150–167 (cit. on pp. 13, 152).

[KS74] G. M. Kelly and R. Street. 'Review of the Elements of 2-Categories'. In: *Category Seminar: Proceedings Sydney Category Seminar 1972/1973*. Ed. by Gregory M. Kelly. Lecture Notes in Mathematics 420. Springer-Verlag, 1974 (cit. on p. 42).

[KKW14] Henning Kerstan, Barbara König and Bram Westerbaan. 'Lifting Adjunctions to Coalgebras to (Re)Discover Automata Constructions'. In: *CMCS*. Vol. 8446. Lecture Notes in Computer Science. Springer, 2014, pp. 168–188 (cit. on p. 65).

[Kle52] Stephen C. Kleene. *Introduction to Metamathematics*. Princeton, New Jersey: D. van Nostrand, 1952 (cit. on p. 120).

[Kli05] Bartek Klin. 'The Least Fibred Lifting and the Expressivity of Coalgebraic Modal Logic'. In: *CALCO*. Vol. 3629. Lecture Notes in Computer Science. Springer, 2005, pp. 247–262 (cit. on pp. 65, 69, 93).

[Kli07] Bartek Klin. 'Coalgebraic Modal Logic Beyond Sets'. In: *MFPS*. Vol. 173. Electronic Notes in Theoretical Computer Science. Elsevier, 2007, pp. 177–201 (cit. on pp. 36, 65, 69, 76, 86).

[Kli11] Bartek Klin. 'Bialgebras for structural operational semantics: An introduction'. In: *Theor. Comput. Sci.* 412.38 (2011), pp. 5043–5069 (cit. on p. 33).

[KS18] Bartek Klin and Julian Salamanca. 'Iterated Covariant Powerset is not a Monad'. In: *MFPS*. Vol. 341. Electronic Notes in Theoretical Computer Science. Elsevier, 2018, pp. 261–276 (cit. on p. 34).

[Kom+19] Yuichi Komorida, Shin-ya Katsumata, Nick Hu, Bartek Klin and Ichiro Hasuo. 'Codensity Games for Bisimilarity'. In: *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, BC, Canada, June 24-27, 2019*. IEEE, 2019, pp. 1–13 (cit. on pp. 27, 29).

[Kom+22] Yuichi Komorida, Shin-ya Katsumata, Nick Hu, Bartek Klin, Samuel Humeau, Clovis Eberhart and Ichiro Hasuo. 'Codensity Games for Bisimilarity'. In: *New Gener. Comput.* 40.2 (2022), pp. 403–465 (cit. on pp. 117, 144, 151, 152).

[Kom+21] Yuichi Komorida, Shin-ya Katsumata, Clemens Kupke, Jurriaan Rot and Ichiro Hasuo. 'Expressivity of Quantitative Modal Logics: Categorical Foundations via Codensity and Approximation'. In: *LICS*. IEEE, 2021, pp. 1–14 (cit. on pp. 6, 12, 66, 117, 144, 151–153).

[KM18] Barbara König and Christina Mika-Michalski. '(Metric) Bisimulation Games and Real-Valued Modal Logics for Coalgebras'. In: *CONCUR*. Vol. 118. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018, 37:1–37:17 (cit. on pp. 12, 117, 144).

[KMS20] Barbara König, Christina Mika-Michalski and Lutz Schröder. 'Explaining Non-bisimilarity in a Coalgebraic Approach: Games and Distinguishing Formulas'. In: *CMCS*. Vol. 12094. Lecture Notes in Computer Science. Springer, 2020, pp. 133–154 (cit. on pp. 13, 117, 123, 145).

[Kor+24] Mayuko Kori, Kazuki Watanabe, Jurriaan Rot and Shin-ya Katsumata. 'Composing Codensity Bisimulations'. In: *LICS*. ACM, 2024, 52:1–52:13 (cit. on p. 151).

[Kor91] Henri Korver. 'Computing Distinguishing Formulas for Branching Bisimulation'. In: *CAV*. Vol. 575. Lecture Notes in Computer Science. Springer, 1991, pp. 13–23 (cit. on p. 13).

[Koz85] Dexter Kozen. 'A Probabilistic PDL'. In: *J. Comput. Syst. Sci.* 30.2 (1985), pp. 162–178 (cit. on p. 134).

[KL51] S. Kullback and R. A. Leibler. 'On Information and Sufficiency'. In: *The Annals of Mathematical Statistics* 22.1 (1951), pp. 79–86 (cit. on p. 144).

[KKP04] Clemens Kupke, Alexander Kurz and Dirk Pattinson. 'Algebraic Semantics for Coalgebraic Logics'. In: *CMCS*. Vol. 106. Electronic Notes in Theoretical Computer Science. Elsevier, 2004, pp. 219–241 (cit. on pp. 36, 65, 66).

[KKP05] Clemens Kupke, Alexander Kurz and Dirk Pattinson. 'Ultrafilter Extensions for Coalgebras'. In: *CALCO*. Vol. 3629. Lecture Notes in Computer Science. Springer, 2005, pp. 263–277 (cit. on pp. 7, 35, 36, 44, 69, 76, 154).

[KR21] Clemens Kupke and Jurriaan Rot. 'Expressive Logics for Coinductive Predicates'. In: *Log. Methods Comput. Sci.* 17.4 (2021) (cit. on pp. 12, 66, 69, 93).

[Kur00] Alexander Kurz. 'Logics for coalgebras and applications to computer science'. PhD thesis. LMU München, 2000 (cit. on p. 34).

[LS91] Kim Guldstrand Larsen and Arne Skou. 'Bisimulation through Probabilistic Testing'. In: *Inf. Comput.* 94.1 (1991), pp. 1–28 (cit. on pp. 10, 98).

[Lei04] T. Leinster. *Higher Operads, Higher Categories*. Vol. 298. London Mathematical Society Lecture Notes. Cambridge University Press, 2004. ISBN: 9781107362277 (cit. on p. 42).

[Lev15] Paul Blain Levy. 'Final Coalgebras from Corecursive Algebras'. In: *CALCO*. Vol. 35. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015, pp. 221–237 (cit. on p. 36).

[Luc+09] Dorel Lucanu, Eugen-Ioan Goriac, Georgiana Caltais and Grigore Rosu. 'CIRC: A Behavioral Verification Tool Based on Circular Coinduction'. In: *CALCO*. Vol. 5728. Lecture Notes in Computer Science. Springer, 2009, pp. 433–442 (cit. on pp. 12, 98).

[LR07] Dorel Lucanu and Grigore Rosu. 'CIRC : A Circular Coinductive Prover'. In: *CALCO*. Vol. 4624. Lecture Notes in Computer Science. Springer, 2007, pp. 372–378 (cit. on p. 12).

[Man69] Ernest Manes. 'A triple theoretic construction of compact algebras'. In: *Seminar on triples and categorical homology theory*. Springer. 1969, pp. 91–118 (cit. on pp. 20, 37).

[Man76] Ernest G Manes. *Algebraic Theories*. Springer New York, 1976 (cit. on p. 20).

[MPP16] Radu Mardare, Prakash Panangaden and Gordon D. Plotkin. 'Quantitative Algebraic Reasoning'. In: *LICS*. ACM, 2016, pp. 700–709 (cit. on pp. 123, 128, 153).

[MG23] Jan Martens and Jan Friso Groote. 'Computing Minimal Distinguishing Hennessy-Milner Formulas is NP-Hard, but Variants are Tractable'. In: *CONCUR*. Vol. 279. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, 32:1–32:17 (cit. on pp. 123, 145).

[MG24] Jan Martens and Jan Friso Groote. 'Minimal Depth Distinguishing Formulas Without Until for Branching Bisimulation'. In: *Logics and Type Systems in Theory and Practice*. Vol. 14560. Lecture Notes in Computer Science. Springer, 2024, pp. 188–202 (cit. on pp. 123, 145).

[Mil80] Robin Milner. *A Calculus of Communicating Systems*. Vol. 92. Lecture Notes in Computer Science. Springer, 1980 (cit. on p. 5).

[Mil89] Robin Milner. *Communication and concurrency*. PHI Series in computer science. Prentice Hall, 1989 (cit. on p. 5).

[Mos99] Lawrence S. Moss. 'Coalgebraic Logic'. In: *Ann. Pure Appl. Log.* 96.1-3 (1999), pp. 277–317 (cit. on pp. 9, 86).

[OHe20] Peter W. O'Hearn. 'Incorrectness logic'. In: *Proc. ACM Program. Lang.* 4.POPL (2020), 10:1–10:32 (cit. on p. 154).

[PT87] Robert Paige and Robert Endre Tarjan. 'Three Partition Refinement Algorithms'. In: *SIAM J. Comput.* 16.6 (1987), pp. 973–989 (cit. on p. 13).

[Par81] David Michael Ritchie Park. 'Concurrency and Automata on Infinite Sequences'. In: *Theoretical Computer Science*. Vol. 104. Lecture Notes in Computer Science. Springer, 1981, pp. 167–183 (cit. on pp. 5, 6).

[Pat01] Dirk Pattinson. 'Expressivity results in the modal logic of coalgebras'. PhD thesis. Universität München, 2001 (cit. on p. 12).

[Pat04] Dirk Pattinson. 'Expressive Logics for Coalgebras via Terminal Sequence Induction'. In: *Notre Dame J. Formal Log.* 45.1 (2004), pp. 19–33 (cit. on p. 12).

[PMW06] Dusko Pavlovic, Michael W. Mislove and James Worrell. 'Testing Semantics: Connecting Processes and Process Logics'. In: *AMAST*. Vol. 4019. Lecture Notes in Computer Science. Springer, 2006, pp. 308–322 (cit. on pp. 36, 65, 76).

[PRT22] Damien Pous, Jurriaan Rot and Ruben Turkenburg. 'Corecursion Up-to via Causal Transformations'. In: *CMCS*. Vol. 13225. Lecture Notes in Computer Science. Springer, 2022, pp. 133–154 (cit. on p. 15).

[Pro56] Yu V Prokhorov. 'Convergence of random processes and limit theorems in probability theory'. In: *Theory of Probability & Its Applications* 1.2 (1956), pp. 157–214 (cit. on p. 144).

[Raa+20] Azalea Raad, Josh Berdine, Hoang-Hai Dang, Derek Dreyer, Peter W. O'Hearn and Jules Villard. 'Local Reasoning About the Presence of Bugs: Incorrectness Separation Logic'. In: *CAV (2)*. Vol. 12225. Lecture Notes in Computer Science. Springer, 2020, pp. 225–252 (cit. on p. 154).

[RB23] Amgad Rady and Franck van Breugel. 'Explainability of Probabilistic Bisimilarity Distances for Labelled Markov Chains'. In: *FoSSaCS*. Vol. 13992. Lecture Notes in Computer Science. Springer, 2023, pp. 285–307 (cit. on pp. 11, 121, 124, 128, 134, 136, 139, 141, 142, 144, 145, 150).

[RG00] Grigore Rosu and Joseph Goguen. 'Circular coinduction'. In: (2000) (cit. on p. 12).

[RL09] Grigore Rosu and Dorel Lucanu. 'Circular Coinduction: A Proof Theoretical Foundation'. In: *CALCO*. Vol. 5728. Lecture Notes in Computer Science. Springer, 2009, pp. 127–144 (cit. on p. 12).

[RJL21] Jurriaan Rot, Bart Jacobs and Paul Blain Levy. 'Steps and traces'. In: *J. Log. Comput.* 31.6 (2021), pp. 1482–1525 (cit. on pp. 8, 35, 36, 42, 44–46, 65, 66, 75, 76).

[RJB24] Jurriaan Rot, Sebastian Junges and Harsh Beohar. 'Relating Apartness and Branching Bisimulation Games'. In: *Logics and Type Systems in Theory and Practice*. Vol. 14560. Lecture Notes in Computer Science. Springer, 2024, pp. 203–213 (cit. on pp. 13, 152).

[Roz24] Wojciech Rozowski. 'A Complete Quantitative Axiomatisation of Behavioural Distance of Regular Expressions'. In: *ICALP*. Vol. 297. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024, 149:1–149:20 (cit. on p. 123).

[Rut98] Jan J. M. M. Rutten. 'Relators and Metric Bisimulations'. In: *CMCS*. Vol. 11. Electronic Notes in Theoretical Computer Science. Elsevier, 1998, pp. 252–258 (cit. on p. 122).

[Rut00] Jan J. M. M. Rutten. 'Universal coalgebra: a theory of systems'. In: *Theor. Comput. Sci.* 249.1 (2000), pp. 3–80 (cit. on pp. 4, 55).

[Sam+21] Michael Sammler, Rodolphe Lepigre, Robbert Krebbers, Kayvan Memarian, Derek Dreyer and Deepak Garg. 'RefinedC: automating the foundational verification of C code with refined ownership types'. In: *PLDI*. ACM, 2021, pp. 158–174 (cit. on p. 154).

[San98] Davide Sangiorgi. 'On the bisimulation proof method'. In: *Mathematical Structures in Computer Science* 8.5 (1998), pp. 447–479 (cit. on p. 6).

[San11] Davide Sangiorgi. *Introduction to bisimulation and coinduction.* Cambridge University Press, 2011 (cit. on p. 6).

[Sch+24] Ezra Schoen, Clemens Kupke, Jurriaan Rot and Ruben Turkenburg. 'A Categorical Approach to Coalgebraic Fixpoint Logic'. In: *CMCS*. Vol. 14617. Lecture Notes in Computer Science. Springer, 2024, pp. 23–43 (cit. on p. 15).

[Sch08] Lutz Schröder. 'Expressivity of coalgebraic modal logic: The limits and beyond'. In: *Theor. Comput. Sci.* 390.2-3 (2008), pp. 230–247 (cit. on pp. 12, 117).

[Sil+13] Alexandra Silva, Filippo Bonchi, Marcello M. Bonsangue and Jan J. M. M. Rutten. 'Generalizing determinization from automata to coalgebras'. In: *Log. Methods Comput. Sci.* 9.1 (2013) (cit. on pp. 7, 33, 65, 83).

[Sok05] Ana Sokolova. 'Coalgebraic analysis of probabilistic systems'. PhD thesis. 2005 (cit. on p. 98).

[Sok11] Ana Sokolova. 'Probabilistic systems coalgebraically: A survey'. In: *Theor. Comput. Sci.* 412.38 (2011), pp. 5095–5110 (cit. on pp. 4, 98, 115).

[SW17] Ana Sokolova and Harald Woracek. 'Termination in Convex Sets of Distributions'. In: *CALCO*. Vol. 72. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017, 22:1–22:16 (cit. on p. 41).

[Spo+24] Timm Spork, Christel Baier, Joost-Pieter Katoen, Jakob Piribauer and Tim Quatmann. 'A Spectrum of Approximate Probabilistic Bisimulations'. In: *CONCUR*. Vol. 311. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024, 37:1–37:19 (cit. on p. 145).

[Spr+21] David Sprunger, Shin-ya Katsumata, Jérémy Dubut and Ichiro Hasuo. 'Fibrational bisimulations and quantitative reasoning: Extended version'. In: *J. Log. Comput.* 31.6 (2021), pp. 1526–1559 (cit. on pp. 6, 27, 29, 81, 117, 144, 151).

[SM17] David Sprunger and Lawrence S. Moss. 'Precongruences and Parametrized Coinduction for Logics for Behavioral Equivalence'. In: *CALCO*. Vol. 72. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017, 23:1–23:15 (cit. on p. 98).

[Sta11] Sam Staton. 'Relating coalgebraic notions of bisimulation'. In: *Log. Methods Comput. Sci.* 7.1 (2011) (cit. on pp. 34, 57, 58, 61, 105).

[Str09] Ross Street. 'Weak distributive laws.' In: *Theory and Applications of Categories [electronic only]* 22 (2009), pp. 313–320 (cit. on p. 21).

[Tan18] Qiyi Tang. 'Computing Probabilistic Bisimilarity Distances'. PhD thesis. York University, Toronto, 2018 (cit. on p. 120).

[Tru71] Kathleen Trustrum. *Linear programming*. Springer, 1971 (cit. on p. 130).

[TP97] Daniele Turi and Gordon D. Plotkin. 'Towards a Mathematical Operational Semantics'. In: *LICS*. IEEE Computer Society, 1997, pp. 280–291 (cit. on p. 33).

[Tur+26] Ruben Turkenburg, Harsh Beohar, Franck van Breugel, Clemens Kupke and Jurriaan Rot. 'Constructing Witnesses for Lower Bounds on Behavioural Distances'. In: *CSL*. Vol. 363. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2026, 25:1–25:22 (cit. on p. 14).

[Tur+23a] Ruben Turkenburg, Harsh Beohar, Clemens Kupke and Jurriaan Rot. 'Forward and Backward Steps in a Fibration'. In: *CALCO*. Vol. 270. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, 6:1–6:18 (cit. on p. 14).

[Tur+24] Ruben Turkenburg, Harsh Beohar, Clemens Kupke and Jurriaan Rot. 'Proving Behavioural Apartness'. In: *CMCS*. Vol. 14617. Lecture Notes in Computer Science. Springer, 2024, pp. 156–173 (cit. on p. 14).

[Tur+23b] Ruben Turkenburg, Clemens Kupke, Jurriaan Rot and Ezra Schoen. 'Preservation and Reflection of Bisimilarity via Invertible Steps'. In: *FoSSaCS*. Vol. 13992. Lecture Notes in Computer Science. Springer, 2023, pp. 328–348 (cit. on p. 14).

[Var03] Daniele Varacca. 'Probability, Nondeterminism and Concurrency: Two Denotational Models for Probabilistic Computation'. PhD thesis. University of Aarhus, 2003 (cit. on p. 34).

[Vie22] Leopold Vietoris. 'Bereiche zweiter ordnung'. In: *Monatshefte für Mathematik und Physik* 32 (1922), pp. 258–280 (cit. on p. 38).

[Vil+08] Cédric Villani et al. *Optimal transport: old and new*. Vol. 338. Springer, 2008 (cit. on p. 124).

[VR99] de Vink,Erik P. and Jan J. M. M. Rutten. 'Bisimulation for Probabilistic Transition Systems: A Coalgebraic Approach'. In: *Theor. Comput. Sci.* 221.1-2 (1999), pp. 271–293 (cit. on p. 98).

[Vla+25] Emily Vlasman, Anto Nanah Ji, James Worrell and Franck van Breugel. 'Explainability is a Game for Probabilistic Bisimilarity Distances'. In: *CONCUR*. 2025 (cit. on p. 144).

[WS20] Paul Wild and Lutz Schröder. 'Characteristic Logics for Behavioural Metrics via Fuzzy Lax Extensions'. In: *CONCUR*. Vol. 171. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020, 27:1–27:23 (cit. on pp. 12, 94).

[Win90] Glynn Winskel. 'A Compositional Proof System on a Category of Labelled Transition Systems'. In: *Inf. Comput.* 87.1/2 (1990), pp. 2–57 (cit. on p. 70).

[Wiß+25] Thorsten Wißmann, Bálint Kocsis, Jurriaan Rot and Ruben Turkenburg. 'Trees in Coalgebra from Generalized Reachability ((Co)algebraic pearl)'. In: *CALCO*. Vol. 342. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025, 15:1–15:18 (cit. on p. 15).

[Wiß+20] Thorsten Wißmann, Stefan Milius, Shin-ya Katsumata and Jérémy Dubut. *A Coalgebraic View on Reachability*. 2020. arXiv: 1901.10717 `[math.CT]` (cit. on p. 110).

[WMS22] Thorsten Wißmann, Stefan Milius and Lutz Schröder. 'Quasilinear-time Computation of Generic Modal Witnesses for Behavioural Inequivalence'. In: *Log. Methods Comput. Sci.* 18.4 (2022) (cit. on pp. 13, 117, 123, 145, 154).

[Wol00] Uwe Wolter. 'On Corelations, Cokernels, and Coequations'. In: *CMCS*. Vol. 33. Electronic Notes in Theoretical Computer Science. Elsevier, 2000, pp. 317–336 (cit. on p. 34).

[Wor05] James Worrell. 'On the final sequence of a finitary set functor'. In: *Theor. Comput. Sci.* 338.1-3 (2005), pp. 184–199 (cit. on pp. 106, 107).

[Wyl06] Oswald Wyler. 'Algebraic theories of continuous lattices'. In: *Continuous Lattices: Proceedings of the Conference on Topological and Categorical Aspects of Continuous Lattices (Workshop IV) Held at the University of Bremen, Germany, November 9–11, 1979*. Springer. 2006, pp. 390–413 (cit. on p. 38).

[ZM22] Maaike Zwart and Dan Marsden. 'No-Go Theorems for Distributive Laws'. In: *Log. Methods Comput. Sci.* 18.1 (2022) (cit. on p. 34).

Summary

State-transition systems are extensively used models in computer science. They consist of the states in which a system can exist, and transitions between these states indicating how the state can evolve. The computers which many of us use can be seen as having a state: the data they store, and transitions: the changes they make to these data to perform computations. In this thesis, we aim to better understand two main aspects of these models: how manipulations of systems affect their behaviour; and how to compare the behaviour of systems.

The first part of this thesis concerns a general notion of transformation of systems seen as coalgebras. In Chapter 3, we give new conditions under which transformations leave the observable behaviour of coalgebras unchanged, in the sense that the states which we consider behaviourally equivalent do not change. In doing so, we provide a new approach for comparing the behaviours of a system before and after a transformation has been applied. The results of Chapter 4 extend this approach to more general properties than behavioural equivalences. We apply our results to the existing transformations of determinisation and ultrafilter extensions, as well as to the setting of coalgebraic modal logic, to provide conditions for adequacy and expressiveness.

In the second part of this thesis, we deal more directly with the comparison of system behaviours. Specifically, we are interested in when systems behave differently. In Chapter 5, we give the definition of *behavioural apartness* for coalgebras, and show how it can be proved by an inductive derivation system. Our most important application is to systems whose state changes probabilistically, not yet captured by earlier work in this direction. For such probabilistic systems, another successful approach for comparing behaviours is the use of behavioural distances. Our work in Chapter 6 gives a sound and (approximately) complete derivation system for a notion of apartness for behavioural distances, namely it establishes lower bounds on the distance between states. Finally, we are able to give a constructive correspondence between proofs and a logic used in earlier work as evidence for distances between states, and apply it to a number of examples.

Samenvatting

Toestand-transitiesystemen zijn veelgebruikte modellen in de informatica. Ze bestaan uit toestanden waarin een systeem zich kan bevinden, en transities tussen deze toestanden die aangeven hoe ze ontwikkelen. De computers die wij veelal gebruiken kunnen bijvoorbeeld gezien worden als systemen met toestanden: de data die ze bevatten, en transities: de veranderingen die ze aan deze data maken om berekeningen uit te voeren. Ons doel in dit proefschrift, is om een beter begrip te ontwikkelen van twee aspecten van zulke modellen: het effect van aanpassingen aan systemen op hun gedrag; en het vergelijken van het gedrag van systemen.

Het eerste deel van dit proefschrift betreft transformaties van systemen opgevat als coalgebra's. In Hoofdstuk 3 geven we nieuwe voorwaarden waaronder transformaties het observeerbare gedrag van coalgebra's onveranderd laten, in de zin dat ze de toestanden die we gedragsequivalent vinden niet veranderen. Ondertussen geven we een nieuwe aanpak voor het vergelijken van het gedrag van systemen voor en na het toepassen van transformaties. De resultaten in Hoofdstuk 4 breiden dit uit naar algemenere eigenschappen van coalgebra's dan alleen de equivalenties in gedrag. We passen onze resultaten toe op de bestaande transformaties determinisatie en ultrafilterextensies, evenals in het kader van coalgebraïsche modale logica, om voorwaarden voor adequaatheid en expressiviteit te verkrijgen.

In het tweede deel van dit proefschrift gaan we verder in op het vergelijken van systeemgedrag. Specifiek zijn we geïnteresseerd in wanneer systemen verschillend gedrag vertonen. In Hoofdstuk 5 geven we de definitie van *gedragsonderscheidbaarheid* in coalgebra's en laten zien dat het bewezen kan worden met een inductief afleidingssysteem. Onze belangrijkste toepassing is op systemen wiens transities plaatsvinden met een zekere willekeurigheid, die nog niet werden omvat door eerder werk. Voor zulke probabilistische systemen is een andere belangrijke methode voor het vergelijken van gedrag de gedragsafstand. Het werk in Hoofdstuk 6 geeft een correct en (bij benadering) volledig afleidingssysteem voor een vorm van onderscheidbaarheid voor gedragsafstanden, namelijk ondergrenzen op de afstand tussen toestanden.

Tot slot laten we een constructieve overeenkomst zien tussen bewijzen in dit systeem en logische formules die in eerder werk zijn gebruikt als getuigen voor gedragsafstanden, en passen we dit toe op een aantal voorbeelden.

Index

Research Data Management

This thesis research has been carried out under the research data management policy of the Institute for Computing and Information Science of Radboud University, The Netherlands. [1]

No research data or code has been produced during this PhD research.

[1] ru.nl/en/institute-for-computing-and-information-sciences/research, last accessed 2025-10-14.

Titles in the IPA Dissertation Series since 2023

S. Kochanthara. *A Changing Landscape: On Safety & Open Source in Automated and Connected Driving*. Faculty of Mathematics and Computer Science, TU/e. 2023-01

L.M. Ochoa Venegas. *Break the Code? Breaking Changes and Their Impact on Software Evolution*. Faculty of Mathematics and Computer Science, TU/e. 2023-02

N. Yang. *Logs and models in engineering complex embedded production software systems*. Faculty of Mathematics and Computer Science, TU/e. 2023-03

J. Cao. *An Independent Timing Analysis for Credit-Based Shaping in Ethernet TSN*. Faculty of Mathematics and Computer Science, TU/e. 2023-04

K. Dokter. *Scheduled Protocol Programming*. Faculty of Mathematics and Natural Sciences, UL. 2023-05

J. Smits. *Strategic Language Workbench Improvements*. Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2023-06

A. Arslanagić. *Minimal Structures for Program Analysis and Verification*. Faculty of Science and Engineering, RUG. 2023-07

M.S. Bouwman. *Supporting Railway Standardisation with Formal Verification*. Faculty of Mathematics and Computer Science, TU/e. 2023-08

S.A.M. Lathouwers. *Exploring Annotations for Deductive Verification*. Faculty of Electrical Engineering, Mathematics & Computer Science, UT. 2023-09

J.H. Stoel. *Solving the Bank, Lightweight Specification and Verification Techniques for Enterprise Software*. Faculty of Mathematics and Computer Science, TU/e. 2023-10

D.M. Groenewegen. *WebDSL: Linguistic Abstractions for Web Programming*. Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2023-11

D.R. do Vale. *On Semantical Methods for Higher-Order Complexity Analysis*. Faculty of Science, Mathematics and Computer Science, RU. 2024-01

M.J.G. Olsthoorn. *More Effective Test Case Generation with Multiple Tribes of AI*. Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2024-02

B. van den Heuvel. *Correctly Communicating Software: Distributed, Asynchronous, and Beyond*. Faculty of Science and Engineering, RUG. 2024-03

H.A. Hiep. *New Foundations for Separation Logic*. Faculty of Mathematics and Natural Sciences, UL. 2024-04

C.E. Brandt. *Test Amplification For and With Developers*. Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2024-05

J.I. Hejderup. *Fine-Grained Analysis of Software Supply Chains*. Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2024-06

J. Jacobs. *Guarantees by construction*. Faculty of Science, Mathematics and Computer Science, RU. 2024-07

O. Bunte. *Cracking OIL: A Formal Perspective on an Industrial DSL for Modelling Control Software*. Faculty of Mathematics and Computer Science, TU/e. 2024-08

R.J.A. Erkens. *Automaton-based Techniques for Optimized Term Rewriting*. Faculty of Mathematics and Computer Science, TU/e. 2024-09

J.J.M. Martens. *The Complexity of Bisimilarity by Partition Refinement*. Faculty of Mathematics and Computer Science, TU/e. 2024-10

L.J. Edixhoven. *Expressive Specification and Verification of Choreographies*. Faculty of Science, OU. 2024-11

J.W.N. Paulus. *On the Expressivity of Typed Concurrent Calculi*. Faculty of Science and Engineering, RUG. 2024-12

J. Denkers. *Domain-Specific Languages for Digital Printing Systems*. Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2024-13

L.H. Applis. *Tool-Driven Quality Assurance for Functional Programming and Machine Learning*. Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2024-14

P. Karkhanis. *Driving the Future: Facilitating C-ITS Service Deployment for Connected and Smart Roadways*. Faculty of Mathematics and Computer Science, TU/e. 2024-15

N.W. Cassee. *Sentiment in Software Engineering*. Faculty of Mathematics and Computer Science, TU/e. 2024-16

H. van Antwerpen. *Declarative Name Binding for Type System Specifications*. Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2025-01

I.N. Mulder. *Proof Automation for Fine-Grained Concurrent Separation Logic*. Faculty of Science, Mathematics and Computer Science, RU. 2025-02

T.S. Badings. *Robust Verification of Stochastic Systems: Guarantees in the Presence of Uncertainty*. Faculty of Science, Mathematics and Computer Science, RU. 2025-03

A.M. Mir. *Machine Learning-assisted Software Analysis*. Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2025-04

L.T. Vinkhuijzen. *Data Structures for Quantum Circuit Verification and How To Compare Them*. Faculty of Mathematics and Natural Sciences, UL. 2025-05

D. van der Wal. *What is the Point? Single-Input-Change Testing a EULYNX Controller*. Faculty of Electrical Engineering, Mathematics & Computer Science, UT. 2025-06

A. Rosset. *Uniform Monad Presentations and Graph Quasitoposes*. Faculty of Sciences, Department of Computer Science, VU. 2025-07

L. Guo. *Higher-Order Termination with Logical Constraints*. Faculty of Science, Mathematics and Computer Science, RU. 2025-08

D.A. Manrique Negrin. *A Model Orchestra in Digital Twins: A Model-Driven Approach to Integration and Orchestration*. Faculty of Mathematics and Computer Science, TU/e. 2025-09

C.A. Esterhuyse. *Specification-Centric Multi-Agent Systems*. Faculty of Science, UvA. 2025-10

H.M. Muctadir. *Consistency Matters: Building Consistent Digital Twin Virtual Entities*. Faculty of Mathematics and Computer Science, TU/e. 2025-11

A. Stramaglia. *Model Checking Machine-Control Applications*. Faculty of Mathematics and Computer Science, TU/e. 2025-12

M. Saeedi Nikoo. *Supporting business process management: clone detection and recommendation techniques*. Faculty of Mathematics and Computer Science, TU/e. 2025-13

R.B. Rubbens. *Bridging the Implementation Gap: Advances in Model-Based Concurrent Program Verification*. Faculty of Electrical Engineering, Mathematics & Computer Science, UT. 2025-14

F.I. van der Berg. *DMC Model Checker: Delta-Driven Variable-Length Next-State Generation via Recursive Compression*. Faculty of Electrical Engineering, Mathematics & Computer Science, UT. 2025-15

D.G. Sprokholt. *Correct Translation between Weak Memory Model Architectures*. Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2025-16

B. Kohlen. *Verified & Fast Probabilistic Verification: Stepwise Refinement of Probabilistic Model Checking into an IEEE 754 Floating-Point Implementation*. Faculty of Electrical Engineering, Mathematics & Computer Science, UT. 2026-01

D. Frolich. *Exploratory Language Development*. Faculty of Science, UvA. 2026-02

L.B. van den Haak. *Verifying Optimised Parallel Code*. Faculty of Mathematics and Computer Science, TU/e. 2026-03

A. Papotti. *Human Acceptance of Machine Generated Code Security Analysis*. Faculty of Sciences, Department of Computer Science, VU. 2026-04

N. Doorn. *O bug, where art thou? Improving software testing education*. Faculty of Science, OU. 2026-05

C.R. van der Rest. *Reusable Programming Language Components*. Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2026-06

T.T.P. Franken. *A Parallel Programming Perspective – Designing AuDaLa and AuDaLogic*. Faculty of Mathematics and Computer Science, TU/e. 2026-07

T.J. Smeding. *Reverse-mode Automatic Differentiation for Functional Array Languages*. Faculty of Science, UU. 2026-08

S.A. Rukmono. *Towards Automated Software Architecture Explanation: An AI-enabled method and industrial evaluation*. Faculty of Mathematics and Computer Science, TU/e. 2026-09

R.T.C. Turkenburg. *Comparing and Constructing (Co)Inductive Predicates*. Faculty of Science, Mathematics and Computer Science, RU. 2026-10

9 789465 152448

Radboud Universiteit