Fully Abstract Models for Effectful λ -Calculi via Category-Theoretic Logical Relations*

OHAD KAMMAR, University of Edinburgh, U.K.

SHIN-YA KATSUMATA, National Institute of Informatics, Japan

PHILIP SAVILLE, University of Oxford, U.K.

We present a construction which, under suitable assumptions, takes a model of Moggi's computational λ -calculus with sum types, effect operations and primitives, and yields a model that is adequate and fully abstract. The construction, which uses the theory of fibrations, categorical glueing, $\top\top$ -lifting, and $\top\top$ -closure, takes inspiration from O'Hearn & Riecke's fully abstract model for PCF. Our construction can be applied in the category of sets and functions, as well as the category of diffeological spaces and smooth maps and the category of quasi-Borel spaces, which have been studied as semantics for differentiable and probabilistic programming.

CCS Concepts: • **Theory of computation** → **Denotational semantics**; **Categorical semantics**.

Additional Key Words and Phrases: full abstraction, call-by-value, O'Hearn & Riecke, fibration, monad

ACM Reference Format:

Ohad Kammar, Shin-ya Katsumata, and Philip Saville. 2022. Fully Abstract Models for Effectful λ -Calculi via Category-Theoretic Logical Relations. *Proc. ACM Program. Lang.* 6, POPL, Article 44 (January 2022), 28 pages. <https://doi.org/10.1145/3498705>

1 INTRODUCTION

Two programs are *contextually equivalent* if they evaluate to the same result (with the same effect) in all program contexts. This natural notion of equality makes precise the intuitive 'sameness' a programmer is generally interested in, for example for optimization or compilation. However, the quantification over all program contexts makes establishing contextual equivalence notoriously difficult. A wide variety of techniques have been proposed to mitigate this difficulty: in this work we focus on a particular strand, namely the construction of *fully abstract denotational models*.

A *denotational semantics* assigns a mathematical object (e.g. a set-theoretic function) to each term or program construct. Two terms are *denotationally equal* if they are assigned the same denotation. If any two terms with the same denotation are contextually equivalent, a model is *adequate*; if any two contextually-equivalent terms have the same denotation, it is *fully abstract*. Thus, adequacy and full abstraction roughly correspond to soundness and completeness in logic: in an adequate and fully abstract model, denotational equality completely characterises contextual equivalence. By constructing fully abstract models, one reduces contextual equivalence to denotational equality.

*PS was supported by the Air Force Office of Scientific Research under award number FA9550-21-1-0038. OK and PS were supported by Royal Society University Research Fellow Enhancement Award number RGF/EA/181059. SK was supported by JST ERATO HASUO Metamathematics for Systems Design Project (No. JPMJER1603).

Authors' addresses: Ohad Kammar, School of Informatics, University of Edinburgh, U.K., ohad.kammar@ed.ac.uk; Shin-ya Katsumata, National Institute of Informatics, Japan, s-katsumata@nii.ac.jp; Philip Saville, Department of Computer Science, University of Oxford, U.K., philip.saville@cs.ox.ac.uk.

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2022 Copyright held by the owner/author(s).

2475-1421/2022/1-ART44

<https://doi.org/10.1145/3498705>

In this paper we construct fully abstract models for a wide class of languages with *effects* encoded by monads (à la Moggi [1989, 1991]), including set-theoretic state (Sec. 9), measure-theoretic probabilistic programming (Sec. 11.1) and state in differentiable neural network programming (Sec. 11.2). Our main theorem (Thm. 10.2) says the following. Suppose one chooses an effect (e.g. global state) and a *signature* of operations (e.g. operations for reading and writing to memory), as well as a *semantic model* consisting of a category $\mathcal{M}$, a monad, and a denotation for each operation. Then, subject to suitable conditions on $\mathcal{M}$, one can construct an adequate and fully abstract model in which the morphisms are maps in $\mathcal{M}$ preserving certain predicates. Our construction is inspired by that of O’Hearn and Riecke [1995], so we call it the *OHR construction*.

1.1 First Steps towards the OHR Construction

Much of the preceding work on full abstraction has focussed on languages with recursion (see Sec. 1.3), but even set-theoretic models for simple languages can fail to be fully abstract. The next example, which we learned from C. Matache & S. Staton, shows the natural model of read-only state for a global one-bit reference cell is not fully abstract. For now we use notation rather loosely.

Example 1.1. Consider a language with a type `bool` of booleans and a signature containing terms `tt` and `ff` for the booleans, an operation `read` of type $(1 \rightarrow \text{bool})$ reading from the reference cell, and operations for the usual logical operations on booleans. This language has a natural semantic interpretation in the category **Fin** of finite sets and functions. We take **Fin** rather than the large category of **Set** of all sets and functions to avoid size issues later on (see Sec. 8); this does not affect the model. On `bool`, set $\llbracket \text{bool} \rrbracket := 2 = \{0, 1\}$. Programs are interpreted using the *reader monad* $RX := (2 \Rightarrow X)$: a term $(\Gamma \vdash M : \sigma)$ is denoted by a function $\llbracket \Gamma \rrbracket \rightarrow R[\llbracket \sigma \rrbracket]$ whose values are functions from the state of the read-only bit to the end result. Product types are interpreted using the cartesian product of sets, and arrow types by sets of functions: $\llbracket \sigma \rightarrow \tau \rrbracket := (\llbracket \sigma \rrbracket \Rightarrow R[\llbracket \tau \rrbracket])$. In particular, the denotation of a closed program of type σ may be identified with an element of $R[\llbracket \sigma \rrbracket]$.

The idea is that programs are parametrised by the value of the reference cell. For example, $\llbracket \text{read}() \rrbracket = \text{id}_2$: if the reference cell contains i , the `read` operation returns i . Similarly, the program $((\text{read}()) \text{ or } \neg(\text{read}()))$ reading from the state twice, negating one result, and taking the disjunction, is interpreted by const_1 , the constant function at 1: the result is 1 no matter what is stored in the cell.

This defines a semantic model: the category is **Fin**, the monad is R , and operations are interpreted as sketched above. Now consider the following closed terms of type $((1 \rightarrow \text{bool}) \rightarrow \text{bool}) \rightarrow \text{bool}$:

$$M := \lambda f. (f \lambda x. \text{tt}) \text{ or } (f \lambda x. \text{ff}) \quad \text{and} \quad M' := \lambda f. (f \text{ read}) \text{ or } (f (\lambda x. \neg(\text{read}())))) \quad (1)$$

Semantically, M and M' are interpreted by elements of $R((1 \Rightarrow R2) \Rightarrow R2) \Rightarrow R2$. As the state can only be read, and never changed, it is intuitively clear that M and M' are contextually equivalent (we give a proof using our construction in Lemma 9.1). However, $\llbracket M \rrbracket \neq \llbracket M' \rrbracket$: if we define $\kappa : (1 \Rightarrow R2) \rightarrow R2$ by $\kappa(g) = \text{const}_1$ if $g(*) = \text{const}_1$ and $\kappa(g) = \text{const}_0$ otherwise, then for any $i, j \in 2$ one has $\llbracket M \rrbracket(i)(\kappa)(j) = 1$ but $\llbracket M' \rrbracket(i)(\kappa)(j) = 0$.

This example highlights the main obstacle to constructing fully abstract models. Typically, a denotational model contains “counterexamples to contextual equivalence”: morphisms such as κ which can be used to distinguish the denotations of contextually equivalent terms. A classic example is the *parallel-or* function, which is commonly used to show the domains model of PCF is not fully abstract [Plotkin 1977]. To obtain a fully abstract model from a non-fully-abstract one, therefore, we need to remove all the counterexamples to contextual equivalence.

How might we refine Ex. 1.1 to remove κ ? Intuitively, κ cannot correspond to a program because it uses too much information. To compute $\kappa(g)$ one must verify that $g(*)$ returns 1 both when the reference cell contains 0 and when the reference cell contains 1. In other words, κ must know how g

behaves in every possible state. But the state is read-only, so programs cannot do this. This suggests that we should restrict to morphisms which do not use such extra information. One way to do this is to construct a new, refined model in which objects are paired with relations and maps are required to preserve these relations. In Ex. 1.2 we outline such a model, again due to C. Matache & S. Staton. Instead of taking just sets, the objects are sets paired with relations R_0 and R_1 which constrain the behaviour of morphisms when the reference cell contains 0 and 1, respectively.

Example 1.2. Let $\mathbb{L}_{\text{Fin}}$ be the category with objects given by triples $(\underline{X}, R_0, R_1)$, where $\underline{X} \in \mathbf{Fin}$ and $R_0, R_1 \subseteq \underline{X}^2$ are binary relations on $\underline{X}$, and morphisms $f : (\underline{X}, R_0, R_1) \rightarrow (\underline{Y}, S_0, S_1)$ given by maps $f : \underline{X} \rightarrow \underline{Y}$ on the carrier sets which preserve both relations: if $(x, x') \in R_i$ then $(fx, fx') \in S_i$. This category has enough structure to model the calculus: it is cartesian closed and has a (strong) monad $\hat{R}$ defined by $\hat{R}(\underline{X}, R_0, R_1) := (R\underline{X}, \hat{R}(R_0), \hat{R}(R_1))$, where $(h, h') \in \hat{R}(R_i)$ if and only if $(hi, h'i) \in R_i$.

We give a denotational semantics to programs in this category by interpreting bool as $I(\text{bool}) := (2, \{(0, 0), (1, 1)\}, \{(0, 0), (1, 1)\})$; this extends to a semantic interpretation of every type using the cartesian closed structure, and the interpretations of the primitives and operations in $\mathbf{Fin}$ all define morphisms in $\mathbb{L}_{\text{Fin}}$ with respect to this new structure. Then one obtains an interpretation in $\mathbb{L}_{\text{Fin}}$ of every term using the cartesian closed structure and the monad $\hat{R}$. We denote the interpretation of $(\Gamma \vdash M : \sigma)$ in $\mathbb{L}_{\text{Fin}}$ by $I[\Gamma \vdash M : \sigma]$ to distinguish it from the interpretation in $\mathbf{Fin}$.

Now, κ is not a morphism $I[1 \rightarrow \text{bool}] \rightarrow \hat{R}(I[\text{bool}])$ in $\mathbb{L}_{\text{Fin}}$. Writing R_0^σ and R_1^σ for the two relations in $I[\sigma]$, then $(\lambda x. \text{id}_2, \lambda x. \text{const}_1) \in R_1^{1 \rightarrow \text{bool}}$ but $(\kappa(\lambda x. \text{id}_2), \kappa(\lambda x. \text{const}_1)) = (\text{const}_0, \text{const}_1)$ is not in R_1^{bool} . In fact, this is sufficient to show κ is not the denotation of any term in the $\mathbf{Fin}$ -model of Ex. 1.1. To see this, note the forgetful functor $U : \mathbb{L}_{\text{Fin}} \rightarrow \mathbf{Fin} : (\underline{X}, R_0, R_1) \mapsto \underline{X}$ strictly preserves the semantic interpretation: $U(I[\Gamma \vdash M : \sigma]) = [\Gamma \vdash M : \sigma]$ for every term. If κ were definable, so that $\kappa = [x : 1 \rightarrow \text{bool} \vdash K : \text{bool}]$, we would get $\kappa = [x : 1 \rightarrow \text{bool} \vdash K : \text{bool}] = U(I[x : 1 \rightarrow \text{bool} \vdash K : \text{bool}])$. Since $U(f) = f$, this entails κ is a map in $\mathbb{L}_{\text{Fin}}$, a contradiction.

The model just sketched is an improvement on that in Ex. 1.1: we have removed the counterexample to contextual equivalence κ . However, this success is only partial. Even though κ is not a morphism in $\mathbb{L}_{\text{Fin}}$, it is still an element of the function space. Since $\kappa \in [(1 \rightarrow \text{bool}) \rightarrow \text{bool}]$ in $\mathbf{Fin}$ and the forgetful functor preserves the semantic interpretation, $\kappa \in U(I[(1 \rightarrow \text{bool}) \rightarrow \text{bool}])$. But then $U(I[M])(i)(\kappa) = [M](i)(\kappa) \neq [M'](i)(\kappa) = U(I[M'])(i)(\kappa)$. Since U is faithful, it follows that $I[M] \neq I[M']$: even though we've cut κ out of our semantic model, its existence in the function space is sufficient to distinguish the denotations of contextually equivalent terms. We solve this using the notion of *concreteness* (cf. O'Hearn and Riecke [1995]).

1.1.1 Cutting Down the Function Space: Concreteness. We start by expressing the property “ κ is not a morphism in $\mathbb{L}_{\text{Fin}}$ but it is an element of the function space” in categorical terms.

Lemma 1.1. *The set map $\ulcorner \kappa \urcorner : 1 \rightarrow [(1 \rightarrow \text{bool}) \rightarrow \text{bool}] : * \mapsto \kappa$ does not define a morphism from the terminal object to $I[(1 \rightarrow \text{bool}) \rightarrow \text{bool}]$ in $\mathbb{L}_{\text{Fin}}$.*

This lemma suggests that we need to restrict our function spaces to consist only of those elements that are ‘named’ by a morphism from the terminal object (a *global element*) in $\mathbb{L}_{\text{Fin}}$. We do this by restricting our attention to the subcategory of *concrete* objects.

Definition 1.1. An object $X := (\underline{X}, R_0, R_1) \in \mathbb{L}_{\text{Fin}}$ is *concrete* if for every $x \in \underline{X}$ the corresponding global element $\ulcorner x \urcorner : 1 \rightarrow \underline{X}$ in $\mathbf{Fin}$ lifts to a global element $\ulcorner x \urcorner : 1 \rightarrow X$ in $\mathbb{L}_{\text{Fin}}$.

Explicitly, $(\underline{X}, R_0, R_1)$ is concrete if for every $x \in \underline{X}$ the pair (x, x) is in both R_0 and R_1 . Write $\mathbb{C}_{\text{Fin}}$ for the full subcategory of $\mathbb{L}_{\text{Fin}}$ consisting of just the concrete objects. This category has enough structure to be a semantic model because it is a *reflective and coreflective subcategory* of $\mathbb{L}_{\text{Fin}}$: the

inclusion $j : \mathbb{C}_{\text{Fin}} \hookrightarrow \mathbb{L}_{\text{Fin}}$ has both left and right adjoints. The left adjoint K adds the diagonal:

$$K(\underline{X}, R_0, R_1) := (\underline{X}, R_0 \cup \{(x, x) \mid x \in \underline{X}\}, R_1 \cup \{(x, x) \mid x \in \underline{X}\}) \quad (2)$$

The right adjoint H , by contrast, cuts down the carrier set. It takes $(\underline{X}, R_0, R_1)$ to the object with carrier $\{x \in \underline{X} \mid (x, x) \in R_0 \text{ and } (x, x) \in R_1\}$ and relations given by restricting R_0 and R_1 to this set.

We can now give the structure for the semantic model. The interpretation of `bool` in $\mathbb{L}_{\text{Fin}}$ is concrete, so the interpretation of base types and constants in $\mathbb{L}_{\text{Fin}}$ restricts to $\mathbb{C}_{\text{Fin}}$. Next, the monad $\hat{R}$ induces a strong monad on $\mathbb{C}_{\text{Fin}}$ with underlying functor $\hat{R}j$ via the adjunction $j \dashv H$. Finally, by the general theory of (co)reflective subcategories, $\mathbb{C}_{\text{Fin}}$ is a cartesian closed category. Products are computed as in $\mathbb{L}_{\text{Fin}}$ but exponentials are not: the function space $(X \Rightarrow_{\mathbb{C}_{\text{Fin}}} Y)$ in $\mathbb{C}_{\text{Fin}}$ is obtained by applying H to the function space $(jX \Rightarrow_{\mathbb{L}_{\text{Fin}}} jY)$ in $\mathbb{L}_{\text{Fin}}$. Thus, the function space in $\mathbb{C}_{\text{Fin}}$ is a version of that in $\mathbb{L}_{\text{Fin}}$ which has been ‘cut down’ by H . The next result makes this explicit.

Lemma 1.2. *For any $X, Y \in \mathbb{C}_{\text{Fin}}$ the carrier set of $(X \Rightarrow_{\mathbb{C}_{\text{Fin}}} Y)$ is isomorphic to $\mathbb{L}_{\text{Fin}}(jX, jY)$.*

The lemma says that we may identify elements of the function space $(X \Rightarrow_{\mathbb{C}_{\text{Fin}}} Y)$ with morphisms in $\mathbb{L}_{\text{Fin}}$. In particular, unwinding the isomorphism shows that, since κ is not a morphism in $\mathbb{L}_{\text{Fin}}$, it cannot be an element of the function space in $\mathbb{C}_{\text{Fin}}$. This also explains why the forgetful functor $U : \mathbb{C}_{\text{Fin}} \rightarrow \mathbf{Fin}$ cannot preserve exponentials—and hence the semantic interpretation—even though it does preserve products: in general $\mathbb{L}_{\text{Fin}}(X, Y) \subsetneq \mathbf{Fin}(\underline{X}, \underline{Y}) = (\underline{X} \Rightarrow \underline{Y})$.

Remark 1.1. The requirement that U does not preserve the semantic interpretation is necessary. If $\mathbb{D}$ is a fully abstract semantic model and the functor $F : \mathbb{D} \rightarrow \mathbf{Fin}$ preserves the semantic interpretation, then the denotations in $\mathbb{D}$ of the terms M and M' in (1) must be equal (by full abstraction) but their images under F , namely $\llbracket M \rrbracket$ and $\llbracket M' \rrbracket$, are not (by Ex. 1.2): a contradiction.

One might hope that the semantic model on $\mathbb{C}_{\text{Fin}}$ is fully abstract. However, this is not the case. Although we have cut κ out of both the hom-sets and the function spaces, other counterexamples to contextual equivalence remain: the notion of ‘relation’ employed in $\mathbb{L}_{\text{Fin}}$ is too weak. Nonetheless, the construction of $\mathbb{C}_{\text{Fin}}$ highlights the two sufficient conditions we will rely on to obtain full abstraction in our OHR construction, and provides a template for how to go about ensuring them.

1.1.2 Sufficient Conditions for Full Abstraction. Our construction of $\mathbb{C}_{\text{Fin}}$ had two stages: first, to remove the counterexample to contextual equivalence κ from being a morphism, then to remove it from the function space. Each stage corresponds to a condition we shall require for full abstraction.

First we want to prevent any counterexamples to contextual equivalence from being morphisms. Intuitively, such counterexamples are morphisms that provide information which is not available within the syntax. It is therefore plausible that if every map $\llbracket \Gamma \rrbracket \rightarrow T[\llbracket \sigma \rrbracket]$ is *definable*—that is, the denotation of some term—then no counterexamples to contextual equivalence can exist. A semantic model satisfying this property is called *fully complete* [Abramsky and Jagadeesan 1994].

Next we want to phrase the property achieved by concreteness as a condition on the underlying category. One way to describe the problem caused by Lemma 1.1 is that in $\mathbb{L}_{\text{Fin}}$ we can have $\llbracket M \rrbracket \neq \llbracket M' \rrbracket$ even though $\llbracket M \rrbracket \circ g = \llbracket M' \rrbracket \circ g$ for every global element $g : 1 \rightarrow \llbracket (1 \rightarrow \text{bool}) \rightarrow \text{bool} \rrbracket$. Hence $\llbracket M \rrbracket$ and $\llbracket M' \rrbracket$ may agree on the denotation of every closed term but still not be equal. This is the property we want to outlaw: we require our semantic model to be *well-pointed*, so two maps $f, f' : X \rightarrow Y$ are equal iff $f \circ g = f' \circ g$ for every $g : 1 \rightarrow X$ (cf. Freyd [1972]; Lawvere [2006]). Indeed, the step from $\mathbb{L}_{\text{Fin}}$ to $\mathbb{C}_{\text{Fin}}$ restricts a non-well-pointed category to a well-pointed one.

Even though we have only provided informal motivation, full completeness and well-pointedness are actually sufficient to guarantee full abstraction: see Prop. 3.1.

Fig. 1. The construction of $\mathbb{C}_{\text{Fin}}$ compared to the OHR construction

1.1.3 A Template for the OHR Construction. Even though $\mathbb{C}_{\text{Fin}}$ is not fully abstract, it does highlight the two key steps our OHR construction will take. We therefore finish our exploration of $\mathbb{C}_{\text{Fin}}$ by presenting its construction in abstract terms. At this stage the details of the definitions are not important: instead one should focus on the broad outline, because the OHR construction will have the same shape. For a more detailed, but still high-level, overview see Sec. 2.

$\mathbb{L}_{\text{Fin}}$ may be constructed using *fibrations for logical relations* [Katsumata 2005, 2013]. These build on the observation of Hermida [1993] and Jacobs [1999] that properties of programming languages are naturally described using (Grothendieck) *fibrations*, and axiomatise the structure required to study logical relations in category-theoretic terms. At this stage our main example arises from the *subobject fibration* $\text{cod} : \text{Sub}(\mathbf{Fin}) \rightarrow \mathbf{Fin}$. The objects of $\text{Sub}(\mathbf{Fin})$ are pairs of finite sets (X, A) such that $A \subseteq X$: we think of A as a unary predicate on X . Morphisms $(X, A) \rightarrow (X', A')$ are functions $f : X \rightarrow X'$ which preserve the predicates. The functor cod takes (X, A) to the superset X .

One obtains $\mathbb{L}_{\text{Fin}}$ by *change-of-base* along the functor $\Delta : X \mapsto (X, X) : \mathbf{Fin} \rightarrow \mathbf{Fin} \times \mathbf{Fin}$. This means $\mathbb{L}_{\text{Fin}}$ is the category constructed as the pullback in Figure 1a. The theory of fibrations for logical relations guarantees that $\mathbb{L}_{\text{Fin}}$ is cartesian closed and that U strictly preserves this structure. The construction is completed by restricting to the full subcategory of concrete objects, as shown. Although we have yet to give the relevant definitions—we sketch these in Sec. 2—Figure 1b shows our OHR construction takes the same form. The category $\mathcal{C}$ is constructed by choosing an appropriate fibration for logical relations, applying change-of-base, then restricting to the full subcategory of concrete objects. Choosing $\mathbb{L}$ appropriately gives the OHR model $\text{OHR}(\mathcal{M})$.

1.2 This Paper

We construct a fully abstract model for λ_c^+ , an extension of Moggi’s computational λ -calculus with sum types, primitives, and effect operations. Our main result is that, for any model of λ_c^+ satisfying suitable conditions, the OHR construction over that model exists and is fully abstract (Thm. 10.2). We take an abstract, category-theoretic approach so that our construction is parametric in the input model and works for any choice of monadic effect.

We start without sum types (Secs. 7 and 8), then refine the construction to include them (Sec. 10). As in O’Hearn & Riecke’s work, the maps in $\text{OHR}(\mathcal{M})$ are maps in $\mathcal{M}$ preserving a certain class of relations, the carrier of the function space is a sub-space of the function space in $\mathcal{M}$, and the canonical functor $\text{OHR}(\mathcal{M}) \rightarrow \mathcal{M}$ strictly preserves sums and products (but *not* exponentials, cf. Remark 1.1). Thus, $\mathcal{M}$ and $\text{OHR}(\mathcal{M})$ are tightly connected: we explore this in Sec. 9. Our development makes use of the abstract framework of *fibrations for logical relations* and *logical relations of varying arity* (Sec. 4), as well as an analysis of the definability predicate for λ_c^+ over an arbitrary model (Sec. 5). We also give a bird’s-eye view of the key steps in the construction before diving into the details (Sec. 2).

As applications, we instantiate the construction for three different languages, each over a different category of sets-with-structure (Secs. 9 and 11). We consider the language for read-only state of

Ex. 1.1, an idealised language for probabilistic programming over the category of quasi-Borel spaces [Heunen et al. 2017], and a simple language for differential programming over the category of diffeological spaces [Huot et al. 2020; Iglesias-Zemmour 2013; Souriau 1980]. In Sec. 9 we also provide a conceptual justification for why the counterexample κ of Ex. 1.1 does not give rise to a counterexample in the OHR model. As well as showcasing the structure of the OHR model, and how one can work with it, this helps explain why the construction succeeds.

Technical Contributions. The key message of this paper is that fully abstract models for λ_c^+ can often be constructed mechanistically, without using special features of the model or language. Indeed, our construction applies to any set-theoretic model of λ_c^+ , independently of the monadic effect, so long as every element of the base types is denoted by some term (see Ex. 8.1).

The main technical contributions are as follows. (1) An analysis of the relationship between logical relations, semantic-interpretation preserving functors, and definability. This includes a characterisation of the definable morphisms in an arbitrary model of λ_c^+ , which builds on the work of Katsumata [2008, 2013] (Sec. 5). (2) The introduction of an abstract, fibrational notion of *concreteness*, which serves to ‘cut down’ a function space to just those morphisms preserving suitable properties (Sec. 6). (3) A conceptual, category-theoretic framework for constructing fully abstract models independently of the choice of monadic effect, base types, and primitives (Thm. 10.2). As well as covering a wide range of examples, this also lays the foundations for future work.

1.3 Related Work

Our analysis of definability—and logical relations more generally—builds on Katsumata’s work on sum types (2008) and computational effects (2005). The theory we develop shows how to combine these two approaches, and also extends from the monadic metalanguage over **Set** to the computational λ -calculus over an arbitrary model. As far as we are aware, this is the first characterisation of definability for the computational λ -calculus, with or without sum types. A very different approach is due to Fiore and Simpson [1999], who characterise definability for sum types in the simply-typed λ -calculus using a sheaf condition. The closest work to our own is that of Goubault-Larrecq et al. [2004], who show that a certain logical relation is sound and complete for contextual equivalence for a version of Moggi’s monadic metalanguage with cryptographic primitives and name generation. This builds on previous work [Lasota et al. 2007], which also relates the monadic lifting of Goubault-Larrecq et al. [2008] to logical relations, and hence a form of full abstraction, but only for specific effects and types up to first-order.

Because morphisms in $\text{OHR}(\mathcal{M})$ are maps in $\mathcal{M}$ which preserve certain relations, one can test equality of morphisms in $\text{OHR}(\mathcal{M})$ just as easily as in $\mathcal{M}$: for example, over a set-theoretic model it is just equality of functions. This contrasts with the fully abstract models constructed by quotienting with an equivalence relation generated from the syntax (e.g. [de’ Liguoro 1996; Milner 1977]), where morphisms are equivalence classes of maps in the original model. Indeed, writing $\text{deL}[\![\sigma]\!]$ for the interpretation of a type in a de’Liguoro-style model and $\mathcal{M}[\![\sigma]\!]$ for its interpretation in the original model, one has $\varphi \in \text{deL}[\![\sigma \rightarrow \tau]\!]$ if and only if there exists some $f \in \mathcal{M}[\![\sigma \rightarrow \sigma]\!]$ such that φ is the equivalence class of f . To verify this is well-defined one then needs to prove that φ induces a map of sets, i.e. an element of **Set**($\text{deL}[\![\sigma]\!]$, $\text{deL}[\![\tau]\!]$), and check it has the structure required to be a morphism in the model (e.g. Scott continuity / smoothness / ...). Our construction avoids this extra work: the required properties on morphisms and function spaces are all imposed by the coreflection H and the fact $\text{OHR}(\mathcal{M})$ is a concrete category over $\mathcal{M}$.

Our approach also differs from that taken in games semantics, which has been highly successful in constructing fully abstract models for a variety of languages (e.g. Abramsky et al. [1998, 2000]; Clairambault and de Visme [2020]; Hyland and Ong [2000]). In such models the notion of morphism

is changed to include more intensional information, and one obtains a fully abstract model by identifying suitable morphisms. Games models are often effectively presentable and may be used to give decision procedures for certain fragments (e.g. [Murawski and Tzevelekos 2012]). Our construction, by contrast, is highly non-effective.

We are not aware of any previous fully abstract model for a Moggi-style language with its usual categorical semantics, even for a fixed effect: the closest we know of is the characterisation of contextual equivalence given by Goubault-Larrecq et al. [2004]. Indeed, much of the work in this area has been inspired by PCF. Plotkin [1977] famously showed that, even though the domains model is not fully abstract for PCF, it is fully abstract for PCF extended with parallel-or. This prompted a rich vein of research attempting to classify *sequential* computation, and hence construct a fully abstract model for PCF in domains (see e.g. Fiore et al. [1996]). More recent non-games-based constructions generally take their inspiration from this rich literature and focus on languages with recursion: beyond the original O’Hearn–Riecke model, notable examples include Cartwright et al. [1994]; Ehrhard et al. [2014]; Marz [2000]; Matache et al. [2021]; Riecke and Sandholm [2002]. Apart from the first two works, these models employ a similar basic idea to ours, quantifying over a range of ‘predictions’ which one eventually instantiates to ensure full abstraction. In contrast to our construction, however, all these works construct a category tailored to the language being studied: while they each consider a rich and subtle language, their strategy does not have the range of examples of our λ_c^+ -based approach.

1.4 Future Work

Given the subtleties already evident in models over **Set**, we believe the conditions we require on the starting λ_c^+ -model to be reasonable. But this work is not the end of the story. Further work is required to cover examples such as presheaf models, or domain-theoretic models of recursion.

We see two immediate next steps. First, loosening the requirement that the original model be well-pointed. We expect that, subject to natural assumptions on the monad and its underlying category, the definability predicate and contextual equivalence predicate (cf. Lasota et al. [2007]; Power and Robinson [2000]) should both satisfy logical relations conditions, so that full abstraction at ground types lifts to all higher types. Second, generalising the construction of the hull functor, perhaps using the *comprehension categories* or *subset types* of Jacobs [1993, 1999]. These extensions may enable us to deal with more advanced models, such as presheaf models for local state.

Finally, we would like to extend this work to encompass recursion. First, we would need to show the argument enriches (for example, over $\omega\mathbf{Cpo}$). Second, we would want to incorporate a notion of approximation to exploit the fact that, in a domain-based model, one may approximate the definable elements from below using suitable compactness properties.

2 EXECUTIVE SUMMARY

In this section we give a high-level overview of our OHR construction. As indicated in Sec. 1.1.3, this broadly follows the steps used to construct $\mathbb{C}_{\text{Fin}}$; the construction is summarised in Figure 1b. We assume a language given by choosing a *signature* S of base types, primitives and effectful operations, together with a *semantic model* consisting of a cartesian closed category $\mathcal{M}$, a (strong) monad T , and a semantic interpretation of the base types and constants. The cartesian closed structure and monad then determine a semantic interpretation $\llbracket - \rrbracket$ of all the terms in the language. In this overview we omit sum types as the construction with them is a little more complex.

2.1 Strengthening the Notion of Relation

We saw in Sec. 1.1.2 that in order to construct a fully abstract semantic model it suffices to construct one that is fully complete and well-pointed. As with $\mathbb{C}_{\text{Fin}}$, we shall handle well-pointedness by

restricting to a subcategory of concrete objects. Our main aim, therefore, is to soup up the notion of ‘relation’ in $\mathbb{L}_{\text{Fin}}$ so that every morphism is definable. Our approach is based on two old insights.

Insight 1: Kripke Relations of Varying Arity. Jung and Tiuryn [1993] observed that, because of the variable binding in λ -abstraction, it is highly non-trivial to characterise the definable morphisms in a model of the simply-typed lambda calculus using n -ary relations for a fixed n . They therefore introduced *Kripke relations of varying arity*: a Kripke relation of varying arity on a set $\underline{X}$ consists of a relation $R(\Gamma) \subseteq ([\Gamma] \Rightarrow \underline{X})$ for every context Γ , compatible with variable renaming and weakening. For brevity we call these simply *Kripke relations*.

Kripke relations and relation-preserving morphisms form a category $\text{Krip}_{\mathcal{M}, [-]}$ (we shall give a precise definition momentarily). This may be thought of as a version of $\mathbb{L}_{\text{Fin}}$ in which the binary relations R_0 and R_1 are replaced by a family of relations with arities given by the contexts. In other words, the relations are replaced by a *presheaf* on a *category of contexts* (cf. [Fiore et al. 1999]).

Definition 2.1. A context Γ is a finitely-supported partial map $\text{Var} \rightarrow_{\text{fin}} \text{Ty}$ from a countably infinite set of variables to the set of types; we write $\diamond$ for the empty context and $(x : \sigma) \in \Gamma$ for $\Gamma(x) = \sigma$. Context renamings $\rho : \Gamma \rightarrow \Delta$ are maps $\rho : \text{Dom } \Gamma \rightarrow \text{Dom } \Delta$ that respect the types: if $(x : \sigma) \in \Gamma$ then $(\rho x : \sigma) \in \Delta$. We write Con_s for the category of contexts and context renamings.

The interpretation $[-]$ defines a functor $\text{Con}_s^{\text{op}} \rightarrow \mathcal{M}$: if $\Gamma := (x_i : \sigma_i)_{i=1, \dots, n}$, $\rho : \Gamma \rightarrow \Delta$ and $\rho(x_i) = y_{\rho i}$ for all i then $[\rho] := \langle \pi_{\rho 1}, \dots, \pi_{\rho n} \rangle : [\Delta] \rightarrow [\Gamma]$. Classically, Kripke relations of varying arity are defined with respect to this functor. For our purposes later we give a slightly more general definition, in which $[-]$ is replaced by an arbitrary functor and Con_s^{op} by an arbitrary category.

Definition 2.2. Let $F : \mathbb{A} \rightarrow \mathcal{M}$ be a functor from a small category $\mathbb{A}$. A *Kripke relation of varying arity with respect to F* , which we call simply a *Kripke relation*, on $\underline{X} \in \mathcal{M}$ is a family of predicates $\{R(\Gamma) \subseteq \mathcal{M}(F(\Gamma), \underline{X})\}_{\Gamma \in \mathbb{A}}$ satisfying the *monotonicity* condition: if $h \in R(\Gamma)$ and $\rho : \Delta \rightarrow \Gamma$ is a morphism in $\mathbb{A}$, then $h \circ F\rho \in R(\Delta)$. Kripke relations form a category $\text{Krip}_{\mathcal{M}, F}$ with objects $(\underline{X}, R)$ where $\underline{X} \in \mathcal{M}$ and R is a Kripke relation on $\underline{X}$. A morphism $f : (\underline{X}, R) \rightarrow (\underline{Y}, S)$ is a map $f : \underline{X} \rightarrow \underline{Y}$ in $\mathcal{M}$ that preserves the relation: if $h \in R(\Gamma)$ then $f \circ h \in S(\Gamma)$.

For $F := [-]$ and $\mathcal{M} \subseteq \text{Set}$ the monotonicity condition says that if $\lambda\gamma. x_\gamma \in R(\Gamma) \subseteq ([\Gamma] \Rightarrow \underline{X})$ and $\rho : \Gamma \rightarrow \Delta$ is a context renaming, then $\lambda\delta. x_{[\rho]\delta} \in R(\Delta) \subseteq ([\Delta] \Rightarrow \underline{X})$. In this setting, $f : (\underline{X}, R) \rightarrow (\underline{Y}, S)$ if and only if $\lambda\gamma. x_\gamma \in R(\Gamma)$ implies $\lambda\gamma. f(x_\gamma) \in S(\Gamma)$ for every context Γ .

We now make the analogy with $\mathbb{L}_{\text{Fin}}$ precise by showing how $\text{Krip}_{\mathcal{M}, F}$ can be constructed in the same way as $\mathbb{L}_{\text{Fin}}$ (recall Figure 1a). As before it is not necessary to follow all the details: our aim is to demonstrate that the same construction is at work. The main step is replacing the fibration for logical relations $\text{cod} : \text{Sub}(\text{Fin}) \rightarrow \text{Fin}$ with its indexed counterpart. We therefore replace subsets by *subpresheaves*: a *sub-presheaf* R of $P : \mathbb{A} \rightarrow \text{Set}$ is a family $\{R(\Gamma) \subseteq P(\Gamma)\}_{\Gamma \in \mathbb{A}}$ compatible with P ’s action on morphisms. The $\mathbb{A}$ -parametrised version of the subobject fibration employed in Figure 1a is the subobject fibration $\text{cod} : \text{Sub}(\hat{\mathbb{A}}) \rightarrow \hat{\mathbb{A}}$ taking a sub-presheaf $R \hookrightarrow P$ to P . The category $\text{Krip}_{\mathcal{M}, F}$ then arises as the pullback (change-of-base) along the *nerve functor* $N_F : X \mapsto \mathcal{M}(F(-), X)$, as in Figure 2.

As for $\mathbb{L}_{\text{Fin}}$, it follows from the theory of logical relations of varying arity that $\text{Krip}_{\mathcal{M}, F}$ is a cartesian closed category, and that the forgetful functor U strictly preserves this structure (see Sec. 4.1). Moreover, we can use the $\top\top$ -lifting of Katsumata [2005] to define a monad $\hat{T}$ on $\text{Krip}_{\mathcal{M}, F}$ so that U strictly preserves the monadic structure: $\hat{T}$ is called a *lifting* of T (Def. 4.2). Thus, $\text{Krip}_{\mathcal{M}, F}$ is a generalised category of relations and it has enough

$$\begin{array}{ccc}
 \hat{T} \curvearrowright \text{Krip}_{\mathcal{M}, F} & \xrightarrow{\quad} & \text{Sub}(\hat{\mathbb{A}}) \\
 \downarrow U & \lrcorner & \downarrow \text{cod} \\
 T \curvearrowright \mathcal{M} & \xrightarrow[N_F]{} & \hat{\mathbb{A}}
 \end{array}$$

Fig. 2. The construction of $\text{Krip}_{\mathcal{M}, F}$

structure to interpret the language of interest. The question of which interpretation to take is the subject of our next insight.

Insight 2: Logical Relations. The second insight, due to Plotkin [1973], is that a morphism is definable only if it *satisfies every logical relation*: this is often referred to as the “fundamental lemma” of logical relations. Classically, a logical relation is a family of relations $\{R_\sigma\}_{\sigma \in \text{Ty}}$ indexed by types which is compatible with the type formation rules: for example, one requires that $(h, h') \in R_{\sigma \rightarrow \tau}$ if and only if $(x, x') \in R_\sigma$ implies $(hx, h'x') \in R_\tau$. The definition in our setting is similar, except one must also take account of variable renamings and the monadic effect. For us, a logical relation R consists of a Kripke relation R_σ on $T[\![\sigma]\!]$ for every type σ , compatible with the type formation rules (Def. 5.1). Thus, R is a family of sets indexed by both types and contexts.

Let us make this more precise. To handle the presence of the effect we use the following operation restricting a Kripke relation to *values*, namely those maps that factor through the monadic unit.

Definition 2.3. For $(TX, R) \in \text{Krip}_{\mathcal{M}, F}$, define $R^{\text{val}}(\Gamma) := \{h : F(\Gamma) \rightarrow X \mid \eta_X \circ h \in R(\Gamma)\}$.

Following the strategy for the simply-typed lambda calculus (e.g. Alimohamed [1995]; Ma and Reynolds [1992]; Mitchell and Scedrov [1993]), we express compatibility with type formation by compatibility with the structure of $\text{Krip}_{\mathcal{M}, [-]}$. For a fixed lifting $\hat{T}$ of T , a *logical relation* is a family of Kripke relations $\{R_\sigma\}_{\sigma \in \text{Ty}}$ satisfying equations such as $([\![\sigma]\!], R_\sigma^{\text{val}}) \Rightarrow (T[\![\tau]\!], R_\tau) = ([\![\sigma \rightarrow \tau]\!], R_{\sigma \rightarrow \tau}^{\text{val}})$ and $\hat{T}([\![\sigma]\!], R_\sigma^{\text{val}}) = (T[\![\sigma]\!], R_\sigma)$ in $\text{Krip}_{\mathcal{M}, [-]}$. This captures exactly the required logical relations conditions. For example, if $\mathcal{M} \subseteq \text{Set}$ the exponential $(X, R) \Rightarrow (Y, S)$ is $(X \Rightarrow Y, R \supset S)$, where $\lambda\gamma. h_\gamma \in (R \supset S)(\Gamma)$ if and only if for all renamings $\rho : \Gamma \rightarrow \Delta$ and $\lambda\delta. x_\delta \in R(\Delta)$ one has $\lambda\delta. h_{[\rho]}(\delta)(x_\delta) \in S(\Delta)$. Thus one recovers a renaming-compatible version of the classic condition.

Example 2.1. (1) The crucial example of a logical relation is the *definability predicate* given by $\text{DEF}_\sigma(\Gamma) := \{[\![\Gamma \vdash M : \sigma]\!] \mid M \text{ is derivable}\}$. Then $f \in \text{DEF}_\sigma^{\text{val}}(\Gamma)$ if and only if $\eta_{[\![\sigma]\!] } \circ f$ denotes a term. (2) The forgetful functor $U : \mathbb{L}_{\text{Fin}} \rightarrow \mathbf{Fin}$ induces a logical relation $\mathcal{U}$ defined by $\mathcal{U}_\sigma(\Gamma) := \mathbb{L}_{\text{Fin}}(I[\![\Gamma]\!], \hat{R}I[\![\sigma]\!]) \subseteq \mathbf{Fin}(s[\![\Gamma]\!], \text{Rs}[\![\sigma]\!])$. This phenomenon holds generally: see Prop. 5.1.

A logical relation $R := \{R_\sigma\}_{\sigma \in \text{Ty}}$ determines a semantic interpretation in $\text{Krip}_{\mathcal{M}, [-]}$. One sets $I^R[\![\beta]\!] := ([\![\beta]\!], R_\beta^{\text{val}})$ and observes the equations making R logical entail that $I^R[\![\sigma]\!] := ([\![\sigma]\!], R_\sigma^{\text{val}})$ and $\hat{T}(I^R[\![\sigma]\!]) := (T[\![\sigma]\!], R_\sigma)$ for every type σ . Following Alimohamed [1995], we say $f : [\![\Gamma]\!] \rightarrow T[\![\sigma]\!]$ in $\mathcal{M}$ *satisfies* R if and only if $f : I^R[\![\Gamma]\!] \rightarrow \hat{T}(I^R[\![\sigma]\!])$ in $\text{Krip}_{\mathcal{M}, [-]}$. One then recovers a version of the fundamental lemma (Thm. 5.1) stating that a morphism is definable if and only if it satisfies every logical relation. If f is definable, so $f = [\![M]\!]$, it lifts to the map $I^R[\![M]\!]$ in $\text{Krip}_{\mathcal{M}, [-]}$ because $U : \text{Krip}_{\mathcal{M}, [-]} \rightarrow \mathcal{M}$ preserves the semantic interpretation. Conversely, if f satisfies every logical relation, it satisfies DEF ; since $\text{id}_{[\![\Gamma]\!]} \in \text{DEF}_\Gamma^{\text{val}}(\Gamma)$, this entails that $f \circ \text{id}_{[\![\Gamma]\!]} = f \in \text{DEF}_\sigma(\Gamma)$.

Summarising, we have the following: (1) restricting to definable morphisms is restricting to those satisfying every logical relation; (2) the morphisms satisfying a logical relation R are exactly those that lift to morphisms respecting the corresponding semantic interpretation in $\text{Krip}_{\mathcal{M}, [-]}$.

From the Two Insights to a Fully Abstract Model. The lesson we take from the preceding is that, in order to obtain full completeness, it suffices to adapt the definition of $\text{Krip}_{\mathcal{M}, [-]}$ to build a semantic model $\mathcal{C}$ in which morphisms necessarily satisfy every logical relation over that model. There is a risk of circularity here: to construct $\mathcal{C}$ we need to refer to logical relations over $\mathcal{C}$, but these can only be defined once we have constructed $\mathcal{C}$. O’Hearn & Riecke’s way out of this apparent loop is a kind of impredicativity, which they compare to the use of reducibility candidates for proving strong normalisation of the polymorphic λ -calculus (e.g. Girard [1989]).

We take a similar approach. We construct C so that objects are paired not just with one relation but with a whole indexed family of them. By carefully choosing a monad W , a semantic interpretation $l[-]$, and the indexing set $\mathbb{I}$, we can ensure that for any type σ and logical relation $\mathcal{R}$ over C there is $i_0 \in \mathbb{I}$ so that the relation at index i_0 for $l[\sigma]$ is exactly $\mathcal{R}_\sigma^{\text{val}}$. Morphisms in C will be morphisms in $\mathcal{M}$ which preserve the relations at every $i \in \mathbb{I}$, so it will follow that any C -morphism must satisfy $\mathcal{R}$. Crucially, if we choose $\mathbb{I}$ large enough, we can ensure i_0 exists using only data over $\mathcal{M}$.

2.2 Constructing the Fully Abstract Model

We start by defining a category $\mathcal{K}$ in which objects are paired with a family of relations indexed by $\mathbb{I}$. For now we leave $\mathbb{I}$ as an arbitrary set; once we have seen the properties we require for full completeness, we shall show how to choose it concretely. To this end we suppose for now that for each $i \in \mathbb{I}$ we have a category $\mathbb{A}_i$ and a functor $F_i : \mathbb{A}_i \rightarrow \mathcal{M}$. The objects of $\mathcal{K}$ then consist of an object $\underline{X} \in \mathcal{M}$ together with a Kripke relation $\bar{X}(i)$ for each $i \in \mathbb{I}$, so that $(\underline{X}, \bar{X}(i)) \in \text{Krip}_{\mathcal{M}, F_i}$; thus, each $\bar{X}(i)$ is a family of predicates $\{\bar{X}(i)(\Gamma) \subseteq \mathcal{M}(F_i(\Gamma), \underline{X})\}_{\Gamma \in \mathbb{A}_i}$ compatible with the action of F_i on morphisms. Morphisms $f : (\underline{X}, \bar{X}) \rightarrow (\underline{Y}, \bar{Y})$ in $\mathcal{K}$ are morphisms $f : \underline{X} \rightarrow \underline{Y}$ in $\mathcal{M}$ which preserve every relation: if $h \in \bar{X}(i)(\Gamma)$ then $f \circ h \in \bar{Y}(i)(\Gamma)$.

Once again this category arises from the theory of fibrations of logical relations. Fibrations for logical relations are closed under small products, so we can take $\prod_{i \in \mathbb{I}} \text{cod} : \prod_{i \in \mathbb{I}} \text{Sub}(\mathbb{A}_i) \rightarrow \prod_{i \in \mathbb{I}} \widehat{\mathbb{A}}_i$ and pull back along $\langle N_{F_i} \rangle_{i \in \mathbb{I}}$ to obtain $\mathcal{K}$ as in Figure 3. As with $\text{Krip}_{\mathcal{M}, F}$, this category is cartesian closed: its structure is given component-wise, and we also define a monad $\hat{T}$ on $\mathcal{K}$ component-wise. The forgetful functor $U : \mathcal{K} \rightarrow \mathcal{M}$ then preserves both the cartesian closed and monadic structure.

The category C of our fully abstract model will be the subcategory $j : C \hookrightarrow \mathcal{K}$ consisting of just the concrete objects, namely those $(\underline{X}, \bar{X}) \in \mathcal{K}$ such that every global element $x : 1 \rightarrow \underline{X}$ lifts to a global element $x : 1 \rightarrow X$ in $\mathcal{K}$ (cf. Def. 1.1). This is always a reflective subcategory and, in the examples we consider, also a coreflective subcategory. Thus C is cartesian closed and acquires a monad W with underlying functor $H\hat{T}j$, in the same way as C_{Fin} . This is summarised in Figure 1b. As for C_{Fin} , the function space $(X \Rightarrow_C Y)$ in C has carrier isomorphic to $\mathcal{K}(jX, jY)$ (see Lemma 6.2), so the function space only morphisms preserving the required relations.

It remains to choose $\mathbb{I}$ and the semantic interpretation $l[-]$. Following O'Hearn & Riecke, we choose $\mathbb{I}$ in such a way that it necessarily contains the data we need. To simplify applications, we intentionally choose $\mathbb{I}$ to be as large as possible: the more permissive we are in the choice of $\mathbb{I}$, the more properties we have for reasoning within the OHR model.

It turns out that it suffices to choose $\mathbb{I}$ and $l[-]$ so that for any logical relation $\mathcal{R} = \{\mathcal{R}_\sigma\}_{\sigma \in \text{Ty}}$ over C there exists $i_0 \in \mathbb{I}$ so that the relation at i_0 of the interpretation of a base type β is precisely $\mathcal{R}_\beta^{\text{val}}$, that is, $l[\beta](i_0) = \mathcal{R}_\beta^{\text{val}}$. Let us show why this is the case. If $l[\beta](i_0) = \mathcal{R}_\beta^{\text{val}}$ for every base type β , the cartesian closed structure of C and the monad W yield $l[\sigma](i_0) = \mathcal{R}_\sigma^{\text{val}}$ and $W(l[\sigma])(i_0) = \mathcal{R}_\sigma$ for every type σ (Prop. 7.1). Morphisms in C preserve the relations at every index, so if $f : l[\Gamma] \rightarrow W(l[\Gamma])$ in C then $f : (l[\Gamma], l[\Gamma](i)) \rightarrow (W(l[\Gamma]), l[\Gamma](i))$ in $\text{Krip}_{\mathcal{M}, F_i}$ for every $i \in \mathbb{I}$; instantiating at i_0 yields $f : (l[\Gamma], \mathcal{R}_\Gamma^{\text{val}}) \rightarrow (W(l[\Gamma]), \mathcal{R}_\sigma)$ in $\text{Krip}_{\mathcal{M}, F_{i_0}}$, so f satisfies $\mathcal{R}$.

Now we turn to choosing $\mathbb{I}$. We take it to be a dependent product over enough sets that we can always find the required i_0 (see Sec. 8). An element $i \in \mathbb{I}$ is therefore a tuple $(\mathbb{A}, F, r, \hat{T})$ in which: (1) $\mathbb{A}$ is a category and F is a functor $\mathbb{A} \rightarrow \mathcal{M}$; (2) r is a map assigning a concrete Kripke

$$\begin{array}{ccc}
 \hat{T} \hookrightarrow \mathcal{K} & \xrightarrow{\quad} & \prod_{i \in \mathbb{I}} \text{Sub}(\widehat{\mathbb{A}}_i) \\
 \downarrow U & \lrcorner & \downarrow \prod_{i \in \mathbb{I}} \text{cod} \\
 T \hookrightarrow \mathcal{M} & \xrightarrow{\langle N_{F_i} \rangle_{i \in \mathbb{I}}} & \prod_{i \in \mathbb{I}} \widehat{\mathbb{A}}_i
 \end{array}$$

Fig. 3. The construction of $\mathcal{K}$

relation over F on $\llbracket \beta \rrbracket$ to every base type β , so that $(\llbracket \beta \rrbracket, r(\beta)) \in \text{Krip}_{\mathcal{M}, F}$, together with a suitable interpretation of constants; (3) $\hat{T}$ is a lifting of T to $\text{Krip}_{\mathcal{M}, F}$, so that the forgetful functor preserves the monadic structure. Note that none of this data refers to C , only to $\mathcal{M}$ or data earlier in the tuple.

To see why this definition works, let $\mathcal{R} := \{\mathcal{R}_\sigma(\Gamma) \subseteq C(l[\Gamma], Wl[\sigma])\}_{\Gamma \in \text{Con}, \sigma \in \text{Ty}}$ be a logical relation over C . For technical reasons we need to assume it satisfies certain *compatibility* and properties (e.g. Def. 5.3), but these are mild. We set $i_0 := (\text{Con}_S^{\text{op}}, U \circ l[-], r, \hat{T}^{\mathcal{R}})$, where $l[-]$ is the semantic interpretation in C defined below, U is the forgetful functor $C \rightarrow \mathcal{M}$, $r : \beta \mapsto \mathcal{R}_\beta^{\text{val}}$, and $\hat{T}^{\mathcal{R}}$ is a monad we construct so that $\hat{T}^{\mathcal{R}}(\mathcal{R}_\sigma^{\text{val}})$ is closely related to $\mathcal{R}_\sigma$ (Lemma 7.3).

We verify first that the choice of r and the equality $\overline{l[\sigma]}(i_0) = \mathcal{R}_\sigma^{\text{val}}$ are well-typed. Via the inclusion $C(l[\Gamma], X) \subseteq \mathcal{M}(l[\Gamma], X)$ we get that $\mathcal{R}$ determines a Ty -indexed family of Kripke relations on $\mathcal{M}$ over the functor $U \circ l[-]$ as follows: $\mathcal{R}_\sigma(\Gamma) \subseteq C(l[\Gamma], Wl[\sigma]) \subseteq \mathcal{M}(U(l[\Gamma]), Wl[\sigma])$. Thus, we may identify $\mathcal{R}_\sigma$ with an object $(Wl[\sigma], \mathcal{R}_\sigma) \in \text{Krip}_{\mathcal{M}, U \circ l[-]}$ and $\mathcal{R}_\beta^{\text{val}}$ is a Kripke relation on $l[\beta]$. So long as the carrier $l[\beta]$ equals $\llbracket \beta \rrbracket$, therefore, the equalities are indeed well-typed.

Finally we show how to guarantee that $\overline{l[\beta]}(i_0) = \mathcal{R}_\beta^{\text{val}}$. We do this by choosing the right semantic interpretation. We have just seen that we must set $l[\beta] = \llbracket \beta \rrbracket$ on carriers. On relations, we read off the choice of relation given by the corresponding index: $\overline{l[\beta]}(\mathbb{A}, F, r, \hat{T}) := r(\beta)$. In particular, for i_0 we get that $\overline{l[\beta]}(i_0) = \mathcal{R}_\beta^{\text{val}}$. As outlined above, it follows that every morphism $l[\Gamma] \rightarrow W(l[\sigma])$ in C satisfies $\mathcal{R}$, hence is definable: the semantic model given by C , $l[-]$ and W is fully complete. Moreover, by our restriction to concrete objects, it is well-pointed. Thus, the model is fully abstract.

2.3 Properties of the OHR Model

We finish this summary by noting some properties of C . Morphisms in C are morphisms in $\mathcal{M}$ which satisfy every compatible logical relation over C . To show a $\mathcal{M}$ -morphism f is not a map in C , therefore, it suffices to find a logical relation over C which f does not satisfy. In the presence of effects, doing this directly this can be difficult. We mitigate this in two ways. First, we show that logical relations can be constructed by constructing a semantic model (Prop. 5.1). Often this can be done using an intuitive relational condition, as in the construction of $\mathbb{L}_{\text{Fin}}$ in Ex. 1.1. Second, we indicate how one may induce logical relations on C using certain logical relations on $\mathcal{M}$ (Sec. 9).

The semantic interpretation of ground types in C coincides with that in $\mathcal{M}$. Moreover, in our examples the monad W on C is a sub-monad of the monad T on $\mathcal{M}$: there is a canonical monad morphism giving a monic $c_X : \underline{WX} \hookrightarrow TX$ for every $X \in C$ (Lemma 6.1). Hence, the interpretations of closed ground terms in C are determined by those in $\mathcal{M}$ (Lemma 8.2). This disappears at higher types because the forgetful functor $C \rightarrow \mathcal{M}$ does not preserve function spaces (cf. Remark 1.1).

It follows that C inherits much of the intuition of the original model. For example, in the OHR model over the model $\mathbb{L}_{\text{Fin}}$ of Ex. 1.1, the maps are set maps, the carriers of the function spaces are subsets of those in Fin , and the carrier of WX is a subset of RX . We examine this further in Sec. 9.

Finally, we do not yet incorporate recursion so our model does not use any form of approximation. While O'Hearn & Riecke use a form of compactness to approximate every term using a countable chain, no such facility is available to us. We hope to pursue this line in the future (see Sec. 1.4).

3 THE COMPUTATIONAL λ -CALCULUS λ_c^+

Syntax. Throughout we shall assume a fixed (but arbitrary) choice of λ_c^+ -signature S , consisting of

- (1) A set $\mathbf{B}$ of *base types*, which we extend with a unit type, binary products, an empty type and binary sums to obtain the set of *ground types*: $\mathbf{G} \ni \gamma ::= \beta \in \mathbf{B} \mid 1 \mid \gamma * \gamma \mid 0 \mid \gamma + \gamma$. Including function types yields the set of *simple types*: $\text{Ty}^+ \ni \sigma ::= \beta \in \mathbf{B} \mid 1 \mid \sigma * \sigma \mid 0 \mid \sigma + \sigma \mid \sigma \rightarrow \sigma$.

- (2) A set $\mathbf{E}$ of (algebraic) *effect operations* [Plotkin and Power 2003] and an assignment $\mathbf{E} \rightarrow \mathbf{G} \times \mathbf{G}$ assigning a *parameter type* α and an *arity type* κ to every $\text{op} \in \mathbf{E}$.
- (3) A set $\mathbf{P}$ of *primitives* and an assignment $\mathbf{P} \rightarrow \mathbf{Ty}^+$ of a simple type κ to every $\xi \in \mathbf{P}$. We require that either $\kappa \in \mathbf{G}$ or $\kappa = \sigma_1 * \dots * \sigma_n \rightarrow \gamma$ where $\gamma \in \mathbf{G}$ and each σ_i is a (possibly *thunked*) ground type: either $\sigma_i \in \mathbf{G}$, or $\sigma_i = (1 \rightarrow \gamma_i)$ and $\gamma_i \in \mathbf{G}$.

We write $\text{op} : \alpha \rightsquigarrow \kappa$ to indicate that $\text{op} \in \mathbf{E}$ has parameter type α and arity type κ , and $\xi : \kappa$ to indicate that $\xi \in \mathbf{P}$ is assigned type κ . Although we do not take advantage of it in this paper, we distinguish between effect operations and primitives so that future developments can employ the better metatheory enjoyed by effect operations (e.g. Kammar and Plotkin [2012]; Katsumata [2013]).

Example 3.1. For a set of exceptions E consider the *exception monad* $(-) + E$ together with an effect operation raise_e raising an exception for each $e \in E$. The corresponding handle_e operation is not algebraic ([Plotkin and Power 2003, Example 3]), but one could add it as a primitive.

The typing rules of the type theory $\lambda_c^+(\mathbf{S})$ are the usual rules for the simply-typed lambda calculus with finite products and finite sum types (e.g. Fiore and Simpson [1999]; Scherer [2017]) together with the rules for effect operations and primitives:

$$\frac{\Gamma \vdash M : \alpha}{\Gamma \vdash \text{op } M : \kappa} \quad (\text{op} : \alpha \rightsquigarrow \kappa) \qquad \frac{}{\Gamma \vdash \xi : \kappa} \quad (\xi : \kappa)$$

Semantic Interpretation. A $\lambda_c^+(\mathbf{S})$ -*model* $(\mathcal{M}, T, s)$ consists of

- (1) A category $\mathcal{M}$ with finite products $(\times, 1)$, finite coproducts $(+, 0)$ and exponentials $\Rightarrow$ (that is, a *bi-cartesian closed category* or *bi-CCC*).
- (2) A functor $T : \mathcal{M} \rightarrow \mathcal{M}$ equipped with natural transformations with components $\mu_A : TTA \rightarrow TA$, $\eta_A : A \rightarrow TA$, and $\text{st}_{A,B} : A \times TB \rightarrow T(A \times B)$ satisfying standard axioms [Kock 1972] (that is, a *strong monad* $(T, \mu, \eta, \text{st})$). Where the context is unambiguous we write simply T .
- (3) An interpretation $s : \mathbf{B} \rightarrow \text{Ob}(\mathcal{M})$ of base types. This extends to an interpretation $s[-] : \mathbf{Ty}^+ \rightarrow \text{Ob}(\mathcal{M})$ of all types; on contexts we set $s[\Gamma] := \prod_{(x:\sigma) \in \Gamma} s[\sigma]$.
- (4) An interpretation of effect operations and primitives: a Kleisli arrow $s[\text{op}] : s[\alpha] \rightarrow Ts[\kappa]$ for every $\text{op} : \alpha \rightsquigarrow \kappa$ and a global element $s[\xi] : 1 \rightarrow s[\kappa]$ for every $\xi : \kappa$.

A *morphism of $\lambda_c^+(\mathbf{S})$ -models* $F : (\mathcal{M}, T, s) \rightarrow (\mathcal{M}', T', s')$ is a functor $F : \mathcal{M} \rightarrow \mathcal{M}'$ which strictly preserves all the structure. Thus, F must strictly preserve the cartesian closed structure and satisfy $FT = T'F$, $F\mu_X = \mu'_{FX}$, $F\eta_X = \eta'_{FX}$, $F\text{st}_{X,Y} = \text{st}'_{FX,FY}$ and $Fs = s'$ for all $X \in \mathbf{M}$.

The interpretation s in a $\lambda_c^+(\mathbf{S})$ -model $(\mathcal{M}, T, s)$ extends to an interpretation $s[-]$ of all $\lambda_c^+(\mathbf{S})$ -terms: $s[\Gamma \vdash M : \tau] : s[\Gamma] \rightarrow T(s[\tau])$. The rules are standard (e.g. Fiore and Simpson [1999]; Moggi [1989]) so we omit them. A $\lambda_c^+(\mathbf{S})$ -model morphism strictly preserves these interpretations: if $F : (\mathcal{M}, T, s) \rightarrow (\mathcal{M}', T', s')$ and $(\Gamma \vdash M : \sigma)$ then $F(s[\Gamma \vdash M : \sigma]) = s'[\Gamma \vdash M : \sigma]$.

We shall also consider the development without sum types. Write $\lambda_c(\mathbf{S})$ for the fragment of $\lambda_c^+(\mathbf{S})$ without any sum types, and $\mathbf{Ty}$ for the subset of $\mathbf{Ty}^+$ defined by $\mathbf{Ty} \ni \sigma ::= \beta \in \mathbf{B} \mid 1 \mid \sigma * \sigma \mid \sigma \rightarrow \sigma$. A $\lambda_c(\mathbf{S})$ -*model* consists of a CCC $(\mathcal{M}, \times, 1, \Rightarrow)$ equipped with a strong monad, an interpretation of base types and an interpretation of each operation op and primitive ξ .

We write $f^\#$ for the *Kleisli extension* $\mu_Y \circ Tf : TX \rightarrow TY$ of $f : X \rightarrow TY$.

Full Abstraction and Full Completeness. We streamline the development by using the *semantic* contextual equivalence of Milner [1977]. We say two $\lambda_c^+(\mathbf{S})$ -terms $(\Gamma \vdash M : \sigma)$ and $(\Gamma \vdash M' : \sigma)$ are *contextually equivalent* in a $\lambda_c^+(\mathbf{S})$ -model $(\mathcal{M}, T, s)$, and write $(\Gamma \vdash M \simeq_{\text{ctx}} M' : \sigma)$, if for every closed ground context $C[-]$ one has $s[\diamond \vdash C[M] : \gamma] = s[\diamond \vdash C[M'] : \gamma]$. The model is *fully abstract* if contextual equivalence implies denotational equality: $s[\Gamma \vdash M : \tau] = s[\Gamma \vdash M' : \tau]$ whenever $(\Gamma \vdash M \simeq_{\text{ctx}} M' : \sigma)$. This notion coincides with the syntactic definition in common examples, such

as the domains model of PCF. The converse to full abstraction—called *adequacy*—follows from the compositionality of the interpretation $s[-]$. We now substantiate the claim made in Sec. 1.1.2.

Proposition 3.1. *Any well-pointed, fully complete $\lambda_c^+(\mathcal{S})$ -model is fully abstract.*

The proof echoes the *definable separability* criterion of Curien [2007]; one uses well-pointedness together with the fact that for any term $(\Gamma \vdash M : \tau)$ and environment $(\diamond \vdash G : \prod_{(x:\sigma) \in \Gamma} \sigma)$ there is a context $C[-]$ such that $s[\Gamma \vdash C[M] : \tau] = s[\Gamma \vdash M : \tau]^\# \circ s[\diamond \vdash G : \prod_{(x:\sigma) \in \Gamma} \sigma]$.

4 THE FIBRATIONAL APPROACH TO DEFINABILITY

In this section we introduce the technical machinery referred to in Sec. 2, namely *fibrations for logical relations*, which axiomatise the data required to study logical relations, $\top\top$ -*lifting*, which allows us to construct monads on categories of relations, and $\top\top$ -*closure*, for handling sum types.

4.1 Kripke Relations of Varying Arity and Fibrations for Logical Relations

We assume familiarity with the basics of fibrations: see e.g. Jacobs [1999] or Loregian and Riehl [2020]. For any fibration $p : \mathbb{E} \rightarrow \mathbb{B}$ and product-preserving functor $N : \mathbb{C} \rightarrow \mathbb{B}$, write $U : \text{Krip}(p, N) \rightarrow \mathbb{C}$ for the fibration obtained by *change-of-base* as in the (pullback) diagram in Figure 4.

Objects of $\text{Krip}(p, N)$ are pairs $(\underline{X} \in \mathbb{C}, R \in \mathbb{E})$ such that $N\underline{X} = pR$; morphisms are pairs $(f : \underline{X} \rightarrow \underline{X}', \hat{f} : R \rightarrow R')$ such that $Nf = p(\hat{f})$. We want to study cases in which $\text{Krip}(p, N)$ is a bi-CCC and U strictly preserves this structure. This is captured by the next definition, due to Katsumata [2013].

$$\begin{array}{ccc} \text{Krip}(p, N) & \xrightarrow{\quad} & \mathbb{E} \\ U \downarrow & \lrcorner & \downarrow p \\ \mathbb{C} & \xrightarrow{\quad N \quad} & \mathbb{B} \end{array}$$

Fig. 4. Constructing $\text{Krip}(p, N)$

Definition 4.1. A *fibration for logical relations* over a bi-CCC $\mathbb{B}$ is a partial order bifibration $p : \mathbb{E} \rightarrow \mathbb{B}$ with small fibres and fibrewise small products such that $\mathbb{E}$ is a bi-CCC and p strictly preserves the bi-CCC structure.

Sec. 1.1.3 employed the category $\text{Sub}(\mathbf{Fin})$ of unary *predicates* on $\mathbf{Fin}$; the functor $\text{Sub}(\mathbf{Fin}) \rightarrow \mathbf{Fin} : (A, X) \mapsto X$ is a fibration for logical relations. For more examples see [Katsumata 2013, §6].

Lemma 4.1 ([Katsumata 2013, Prop. 6]). *For any fibration for logical relations $p : \mathbb{E} \rightarrow \mathbb{B}$, bi-CCC $\mathbb{C}$ and product-preserving functor $N : \mathbb{C} \rightarrow \mathbb{B}$, $U : \text{Krip}(p, N) \rightarrow \mathbb{C}$ is a fibration for logical relations.*

As noted above, the categories $\mathbb{L}_{\mathbf{Fin}}$, $\text{Krip}_{\mathcal{M}, F}$ and $\mathcal{K}$ all arise as instances of this lemma. In particular, where N_F denotes the nerve functor $X \mapsto \mathcal{M}(F(-), X)$, we have $\text{Krip}_{\mathcal{M}, F} := \text{Krip}(\text{cod}, N_F)$ and $\mathcal{K} := \text{Krip}(\prod_i \text{cod}, N_{\langle F_i \rangle_i})$. The bi-CCC structure of $\text{Krip}_{\mathcal{M}, F}$ is given below. The exponentials encode the Jung–Tiuryn *logical relations condition* (cf. Jung and Tiuryn [1993]; Plotkin [1980]).

Products: The terminal object is $(1_{\mathbb{C}}, \top)$, where $\top(\Gamma) = \{!_{F\Gamma}\}$. The binary product is $(\underline{A}_1, R_1) \times (\underline{A}_2, R_2) := (\underline{A}_1 \times \underline{A}_2, R_1 \otimes R_2)$, where $h \in (R_1 \otimes R_2)(\Gamma)$ if and only if $\pi_i \circ h \in R_i(\Gamma)$ for $i = 1, 2$.

Exponentials: $(\underline{A}, R) \Rightarrow (\underline{B}, S) := (\underline{A} \Rightarrow \underline{B}, R \supset S)$, where $h \in (R \supset S)(\Gamma)$ if and only if $\text{eval} \circ \langle h \circ F\rho, u \rangle \in S(\Delta)$ for every morphism $\rho : \Delta \rightarrow \Gamma$ and every $u \in R(\Delta)$.

Coproducts: The initial object is $(0_{\mathbb{C}}, \perp)$, where $\perp(\Gamma) = \emptyset$. The binary coproduct is $(\underline{A}_1, R_1) + (\underline{A}_2, R_2) := (\underline{A}_1 + \underline{A}_2, R_1 \oplus R_2)$, with $(R_1 \oplus R_2)(\Gamma) := \{\text{inj}_1 \circ h \mid h \in R_1(\Gamma)\} \cup \{\text{inj}_2 \circ h \mid h \in R_2(\Gamma)\}$.

4.2 Monad Liftings and Generalised $\top\top$ -Lifting

By Lemma 4.1, if we start from a $\lambda_c^+(\mathcal{S})$ -model $(\mathcal{M}, T, s)$ then the category of Kripke relations $\text{Krip}_{\mathcal{M}, s[-]}$ is a bi-CCC. To make this into a $\lambda_c^+(\mathcal{S})$ -model we want a *lifting* of T .

Definition 4.2 ([Katsumata 2005, 2013]). Let $p : \mathbb{E} \rightarrow \mathbb{B}$ be a fibration for logical relations and $(T, \mu, \eta, \text{st})$ a strong monad on $\mathbb{B}$. A *lifting* of T is a strong monad $(\hat{T}, \hat{\mu}, \hat{\eta}, \hat{\text{st}})$ on $\mathbb{E}$ such that $p\hat{T} = Tp$, $p\hat{\mu} = \mu_p$, $p\hat{\eta} = \eta_p$, and $p\hat{\text{st}} = \text{st}_p$.

Our characterisation of definability will rely on the *semantic $\top\top$ -lifting* of [Katsumata \[2005\]](#). Alternative liftings include the *free lifting* [[Kammar and McDermott 2018](#)] and the monadic lifting of [Goubault-Larrecq et al. \[2008\]](#). We choose Katsumata's approach because it interacts well with the extra structure we shall need at sum types (see Def. 5.2).

The next lemma describes the abstract situation.

Lemma 4.2. *Let $p : \mathbb{E} \rightarrow \mathbb{B}$ be a fibration for logical relations and T and S be strong monads on $\mathbb{B}$ related by a morphism of strong monads $\alpha : T \Rightarrow S$ [[Pareigis 1977](#); [Street 1972](#)]. Further assume that S has a lifting $\hat{S}$ and set $\hat{T}X$ to be the cartesian lifting below:*

$$\begin{array}{ccc} \hat{T}X & \xrightarrow{\hat{\alpha}_X} & \hat{S}X \\ T(pX) & \xrightarrow{\alpha_{pX}} & S(pX) \end{array} \quad \begin{array}{c} \mathbb{E} \\ \downarrow p \\ \mathbb{B} \end{array} \quad (3)$$

Then $\hat{T}(-)$ extends to a monad which is a lifting of T , and $\hat{\alpha}$ extends to a morphism of strong monads.

$\top\top$ -lifting arises by taking S and $\hat{S}$ to both be the (strong) continuation monad. In any CCC $\mathbb{C}$, let K_P denote the continuation monad $(- \Rightarrow P) \Rightarrow P$; there exists a canonical monad morphism $\sigma^P : T \Rightarrow K_{TP}$ with components $\lambda(\text{eval}^\# \circ \text{st} \circ \langle \pi_2, \pi_1 \rangle)$. When $p : \mathbb{E} \rightarrow \mathbb{B}$ is a fibration for logical relations, write $\hat{K}_P$ for the continuation monad on $\mathbb{E}$ with target P and observe that this is a lifting of $K_{p(P)}$. We call an object $P \in \mathbb{E}$ such that $p(P) = TQ$ for some $Q \in \mathbb{B}$ a $\top\top$ -lifting parameter.

Definition 4.3 [[Katsumata 2005](#)]. Let $p : \mathbb{E} \rightarrow \mathbb{B}$ be a fibration for logical relations and T be a strong monad on $\mathbb{B}$. Where $p(P) = TQ$ is a $\top\top$ -lifting parameter, the $\top\top[P]$ -lifting $T^{\top\top[P]}$ of T is the monad obtained by taking $\hat{S} := \hat{K}_P$, $S := K_{TQ}$ and $\alpha := \sigma^Q$ in (3). For a small set $\mathbb{P} \neq \emptyset$ of $\top\top$ -lifting parameters, the $\top\top[\mathbb{P}]$ -lifting of T is the fibred product $T^{\top\top[\mathbb{P}]} := \bigwedge_{P \in \mathbb{P}} T^{\top\top[P]}(-)$.

We shall implicitly assume that every set of $\top\top$ -lifting parameters is both non-empty and small.

4.3 Generalised $\top\top$ -Closure

While the definability predicate for a $\lambda_c^+(S)$ -model $(\mathcal{M}, T, s)$ interacts well with the cartesian closed structure of $\text{Krip}_{\mathcal{M}, s}[-]$, this is not the case for the cocartesian structure (see Sec. 5). We therefore follow [Katsumata \[2008\]](#) in restricting to a subcategory with a different coproduct structure.

Definition 4.4. Let $p : \mathbb{E} \rightarrow \mathbb{B}$ be a fibration for logical relations, T be a strong monad on $\mathbb{B}$, and $\hat{T}$ be a lifting of T to $\mathbb{E}$. Define a monad $(-)^{\top\top[\hat{T}]}$ lifting the identity monad by taking $\hat{S} := \hat{T}$, $S := T$, $T := \text{id}$ and $\alpha := \eta^T$ in (3). If $X = X^{\top\top[\hat{T}]}$ we say X is $\hat{T}$ -closed. We write $i : \mathbb{E}^{\top\top[\hat{T}]} \hookrightarrow \mathbb{E}$ for the full subcategory of $\hat{T}$ -closed objects.

The $\top\top$ -closure of [Katsumata \[2008\]](#) arises as a special case. Where $\mathbb{P}$ is a set of $\top\top$ -lifting parameters, an object $X \in \mathbb{E}$ is $\top\top[\mathbb{P}]$ -closed, or simply $\top\top$ -closed, if $\text{id}^{\top\top[\mathbb{P}]}(X) = X$.

Lemma 4.3. *Let $p : \mathbb{E} \rightarrow \mathbb{B}$ be a fibration for logical relations and T be a strong monad on $\mathbb{B}$. For any set $\mathbb{P}$ of $\top\top$ -lifting parameters, $(-)^{\top\top[T^{\top\top[\mathbb{P}}]]} = \text{id}^{\top\top[\mathbb{P}]}$.*

Hence, X is $\top\top[\mathbb{P}]$ -closed if and only if it is $T^{\top\top[\mathbb{P}]}$ -closed. We finish this section by recording some elementary properties of the subcategory of $\hat{T}$ -closed objects. The first two cases generalise properties of the subcategory of $\top\top$ -closed objects proven in [[Katsumata 2008](#)].

Lemma 4.4. *Let $p : \mathbb{E} \rightarrow \mathbb{B}$ be a fibration for logical relations, T be a strong monad on $\mathbb{B}$, and $\hat{T}$ be a lifting of T to $\mathbb{E}$. Then:*

- (1) $(-)^{\top\top[\hat{T}]}$ is an idempotent monad, so $\mathbb{E}^{\top\top[\hat{T}]}$ is a replete, reflective subcategory;
- (2) $\mathbb{E}^{\top\top[\hat{T}]}$ is a sub-CCC of $\mathbb{E}$, and has finite coproducts as $X_i \xrightarrow{\text{inj}_i} \sum_{i \leq n} X_i \xrightarrow{\leq} \text{id}^{\top\top[\hat{T}]}(\sum_{i \leq n} X_i)$;
- (3) $\hat{T}$ restricts to a strong monad on $\mathbb{E}^{\top\top[\hat{T}]}$: if X is $\hat{T}$ -closed, then so is $\hat{T}X$.

5 DEFINABILITY AND $\lambda_c^+(\mathbf{S})$ -LOGICAL RELATIONS

In Sec. 2.1 our second key observation was that the definable morphisms in a fixed $\lambda_c^+(\mathbf{S})$ -model $(\mathcal{M}, T, s)$ may be characterised as those which *satisfy every logical relation*, and that this can be expressed in categorical terms: a map $f : s[\Gamma] \rightarrow T(s[\sigma])$ in $\mathcal{M}$ satisfies R if and only if it lifts to a morphism of Kripke relations $f : l^R[\Gamma] \rightarrow \hat{T}(l^R[\sigma])$. We now make this precise.

Our starting point is the next definition, which extends both logical relations for the simply-typed λ -calculus and logical relations for the monadic metalanguage over **Set** [Katsumata 2005]. Since we use $\lambda_c^+(\mathbf{S})$ rather than the monadic metalanguage we use the $(-)^{\text{val}}$ operation (Def. 2.3).

Definition 5.1. Let $\hat{T}$ be a lifting of T to $\text{Krip}_{\mathcal{M}, s[-]}$. A $\lambda_c(\mathbf{S})$ -logical relation over $\hat{T}$ is a **Ty**-indexed family $\{R_\sigma \hookrightarrow \mathcal{M}(s[-], Ts[\sigma])\}_{\sigma \in \text{Ty}}$ such that for every $\sigma, \tau, \sigma_1, \sigma_2 \in \text{Ty}$ we have

$$R_1^{\text{val}} = \top, \quad R_{\sigma_1 * \sigma_2}^{\text{val}} = R_{\sigma_1}^{\text{val}} \otimes R_{\sigma_2}^{\text{val}}, \quad R_{\sigma \rightarrow \tau}^{\text{val}} = (R_\sigma^{\text{val}} \supset R_\tau), \quad \hat{T}(s[\sigma], R_\sigma^{\text{val}}) = (Ts[\sigma], R_\sigma).$$

Moreover, we require that the interpretation of every effect operation $\text{op} : \alpha \rightsquigarrow \kappa$ and primitive $(\xi : \kappa)$ lifts to Kripke relations: $s[\text{op}] : (s[\alpha], R_\alpha^{\text{val}}) \rightarrow (Ts[\kappa], R_\kappa)$ and $s[\xi] : (1, \top) \rightarrow (s[\kappa], R_\kappa^{\text{val}})$. We extend the definition to contexts using the product structure, so that $R_\Gamma^{\text{val}} := \otimes_{(x:\sigma) \in \Gamma} R_\sigma^{\text{val}}$.

A $\lambda_c(\mathbf{S})$ -logical relation is equivalently an interpretation of base types $\hat{s}^R : \beta \mapsto (s[\beta], R_\beta^{\text{val}})$ which extends to an interpretation satisfying $\hat{s}^R[\sigma] = (s[\sigma], R_\sigma^{\text{val}})$ for every $\sigma \in \text{Ty}$.

We cannot use the coproduct structure of $\text{Krip}_{\mathcal{M}, s[-]}$ to incorporate sum types because elements of $\text{DEF}_{\sigma_1}^{\text{val}} \oplus \text{DEF}_{\sigma_2}^{\text{val}}$ must factor through one of the injections $s[\sigma_i] \rightarrow s[\sigma_1] + s[\sigma_2]$, so one has strict inclusions $\perp \subsetneq \text{DEF}_0^{\text{val}}$ and $(\text{DEF}_{\sigma_1}^{\text{val}} \oplus \text{DEF}_{\sigma_2}^{\text{val}}) \subsetneq \text{DEF}_{\sigma_1 + \sigma_2}^{\text{val}}$. We therefore restrict to the subcategory of $\hat{T}$ -closed objects. Using Lemma 4.4, one obtains the following diagram for any lifting $\hat{T}$ of T . The inclusion i strictly preserves cartesian closed structure but does *not* preserve coproducts.

$$\hat{T} \hookrightarrow \text{Krip}_{\mathcal{M}, s[-]}^{\top\top[\hat{T}]} \xleftarrow[\hat{i}]{(-)^{\top\top[\hat{T}]}} \text{Krip}_{\mathcal{M}, s[-]} \xleftarrow{\hat{T}} \hat{T} \quad (4)$$

We now define $\lambda_c^+(\mathbf{S})$ -logical relations by extending Def. 5.2 with cases for sums; as in Def. 5.2, this extends to contexts using the product structure.

Definition 5.2 (cf. [Katsumata 2008, §3.4]). Let $\hat{T}$ be a lifting of T to $\text{Krip}_{\mathcal{M}, s[-]}$. A $\lambda_c^+(\mathbf{S})$ -logical relation over $\hat{T}$ is a **Ty**⁺-indexed family $\{R_\sigma \hookrightarrow \mathcal{M}(s[-], Ts[\sigma])\}_{\sigma \in \text{Ty}^+}$ such that: (1) the conditions of Def. 5.1 hold; and (2) $R_0^{\text{val}} = \perp^{\top\top[\hat{T}]}$ and $(R_{\sigma_1}^{\text{val}} \oplus R_{\sigma_2}^{\text{val}})^{\top\top[\hat{T}]} = R_{\sigma_1 + \sigma_2}^{\text{val}}$ for every $\sigma_1, \sigma_2 \in \text{Ty}^+$.

Remark 5.1. When $\hat{T}$ is the $\top\top$ -lifting $T^{\top\top[\mathbb{P}]}$, so that $R^{\top\top[\hat{T}]} = \text{id}^{\top\top[\mathbb{P}]}R$, it suffices to require that $((R_{\sigma_1}^{\text{val}} \oplus R_{\sigma_2}^{\text{val}}) \supset Q) = (R_{\sigma_1 + \sigma_2}^{\text{val}} \supset Q)$ and $(R_0^{\text{val}} \supset Q) = (\perp \supset Q)$ for any $Q \in \mathbb{P}$ and $\sigma_1, \sigma_2 \in \text{Ty}^+$: modulo the $(-)^{\text{val}}$ operation, this is the condition in Katsumata's definition.

$\lambda_c^+(\mathbf{S})$ -logical relations satisfy the same key property as $\lambda_c(\mathbf{S})$ -logical relations.

Lemma 5.1. *Let R be a $\lambda_c^+(\mathbf{S})$ -logical relation over $\hat{T}$. Then every $(s[\sigma], R_\sigma^{\text{val}})$ and $(Ts[\sigma], R_\sigma)$ is $\hat{T}$ -closed. Hence if $\hat{s}^R(\beta) = (s[\beta], R_\beta^{\text{val}})$ for all base types β then $\hat{s}^R[\sigma] = (s[\sigma], R_\sigma^{\text{val}})$ for all $\sigma \in \text{Ty}^+$.*

We now characterise the definable morphisms as those which satisfy every $\lambda_c^+(\mathbf{S})$ -logical relation. To do so, we take a detour through $\lambda_c^+(\mathbf{S})$ -model morphisms. We saw in Ex. 1.2 that such functors exclude ‘counterexample’ morphisms from a model; it turns out this approach is implicitly using $\lambda_c^+(\mathbf{S})$ -logical relations. As well as being a key step in our characterisation of definability, the next result is useful because it is often more intuitive to construct a model than a $\lambda_c^+(\mathbf{S})$ -logical relation.

Let $F : (\mathcal{M}, T, s) \rightarrow (\mathcal{N}, S, t)$ be a morphism of $\lambda_c^+(\mathbf{S})$ -models. Since F strictly preserves all the structure, we obtain $\mathcal{F}_\sigma(\Gamma) := \{Fh \mid h \in \mathcal{M}(s[\Gamma], Ts[\sigma])\} \subseteq \mathcal{N}(t[\Gamma], St[\sigma])$ for every $\sigma \in \mathbf{Ty}^+$.

Notation 5.1. Every family $R := \{(Ts[\sigma], R_\sigma)\}_{\sigma \in \mathbf{Ty}^+}$ and hence every $\lambda_c^+(\mathbf{S})$ -logical relation, determines a family of $\top\top$ -lifting parameters: we also denote this by R .

Proposition 5.1. *For any strict $\lambda_c^+(\mathbf{S})$ -model morphism $F : (\mathcal{M}, T, s) \rightarrow (\mathcal{N}, S, t)$, the family $\mathcal{F} := \{\mathcal{F}_\sigma\}_{\sigma \in \mathbf{Ty}^+}$ defined above is a $\lambda_c^+(\mathbf{S})$ -logical relation over the monad $S^{\top\top[\mathcal{F}]}$ lifting S to $\mathbf{Krip}_{\mathcal{N}, t[-]}$.*

Thus, we may recast our construction of the category $\mathbb{L}_{\mathbf{Fin}}$ in Ex. 1.2 as restricting to the subcategory of $\mathbf{Set}$ in which morphisms preserve the $\lambda_c(\mathbf{S})$ -logical relation $\mathcal{U}$ of Ex. 2.1.

Example 5.1. The definability predicate $\mathbf{DEF}$ (Ex. 2.1) arises from Prop. 5.1: for a model $(\mathcal{M}, T, s)$ take the subcategory of $\mathcal{M}$ with objects $\{s[\sigma]\}_{\sigma \in \mathbf{Ty}^+}$ and morphisms the definable maps.

The desired characterisation now follows from Prop. 5.1, Ex. 5.1, and the argument sketched in Sec. 2.1; a similar result holds for $\lambda_c(\mathbf{S})$ -relations. Say that a morphism $f : s[\Gamma] \rightarrow Ts[\sigma]$ in $\mathcal{M}$ satisfies a $\lambda_c^+(\mathbf{S})$ -logical relation $\{R_\sigma\}_{\sigma \in \mathbf{Ty}^+}$ if f lifts to a morphism $\hat{s}^R[\Gamma] \rightarrow \hat{T}(\hat{s}^R[\sigma])$ in $\mathbf{Krip}_{\mathcal{M}, s[-]}^{\top\top[\hat{T}]}$.

Theorem 5.1. $\mathbf{DEF} := \{\mathbf{DEF}_\sigma\}_{\sigma \in \mathbf{Ty}^+}$ is a $\lambda_c^+(\mathbf{S})$ -logical relation over $T^{\top\top[\mathbf{DEF}]}$. Hence, $f : s[\Gamma] \rightarrow Ts[\sigma]$ in $\mathcal{M}$ is $\lambda_c^+(\mathbf{S})$ -definable if and only if it satisfies every $\lambda_c^+(\mathbf{S})$ -logical relation over $T^{\top\top[\mathbf{DEF}]}$.

The logical relations arising via Prop. 5.1 are particularly well-behaved; for example, they are closed under currying. We end this section by axiomatising these properties.

Definition 5.3. A $\lambda_c^+(\mathbf{S})$ -logical relation $R = \{R_\sigma\}_{\sigma \in \mathbf{Ty}^+}$ over $\hat{T}$ is: (1) *hungry* if $\text{id}_{s[\sigma]} \in R_\sigma^{\text{val}}(x : \sigma)$ for all $\sigma \in \mathbf{Ty}^+$; (2) λ -*compatible* if $f \in R_\tau(\Gamma, x : \sigma)$ implies $\lambda f \in R_{\sigma \rightarrow \tau}^{\text{val}}(\Gamma)$; (3) 0 -*compatible* if $h \in R_0(\Gamma)$ implies $T!_{s[\sigma]} \circ h \in R_\tau(\Gamma)$; and (4) $+$ -*compatible* if the composite below is in $R_\tau(\Gamma, p : \sigma_1 + \sigma_2)$ whenever $h_i \in R_\tau(\Gamma, x_i : \sigma_i)$ for $i = 1, 2$ (the isomorphism is the canonical one from distributivity):

$$s[\Gamma, p : \sigma_1 + \sigma_2] = s[\Gamma] \times (s[\sigma_1] + s[\sigma_2]) \xrightarrow{\cong} \sum_{i=1,2} (s[\Gamma] \times s[\sigma_i]) \xrightarrow{[h_1, h_2]} Ts[\tau]$$

If R is λ -compatible, $+$ -compatible and 0 -compatible, we say R is $(\lambda, +, 0)$ -compatible.

A hungry logical relation is one which ‘eats’ every morphism which preserves it: R is hungry if and only if $(f \text{ satisfies } R \implies f \in R_\sigma(\Gamma))$ for every $f : s[\Gamma] \rightarrow Ts[\sigma]$ in $\mathcal{M}$.

6 CONCRETENESS

In Sec. 1.1.1 we obtained a well-pointed model by introducing a criterion which ensures the function spaces only consist of elements that are ‘named’ by a global element, namely *concreteness*. In this section we define concreteness in fibrational terms and show that a $\lambda_c^+(\mathbf{S})$ -model structure on $\mathbf{Krip}(p, \mathbf{N})$ induces a $\lambda_c^+(\mathbf{S})$ -model structure on the subcategory of concrete objects, for a fixed bi-CCC $\mathbb{C}$, fibration for logical relations $p : \mathbb{E} \rightarrow \mathbb{B}$, and product-preserving functor $\mathbf{N} : \mathbb{C} \rightarrow \mathbb{B}$. We shall then substantiate the claim in Sec. 1.1.1 that the function space in the subcategory of concrete objects consists only of those functions which preserve the relevant relations (Lemma 6.2).

Definition 6.1 (cf. Def. 1.1). An object $(\underline{X}, R) \in \mathbf{Krip}(p, \mathbf{N})$ is *concrete* if every global element $g : 1 \rightarrow \underline{X}$ in $\mathbb{C}$ lifts to a global element $(g, \hat{g}) : (1, \top) \rightarrow (\underline{X}, R)$ lying over g . We write

$j : \text{Conc}(p, N) \hookrightarrow \text{Krip}(p, N)$ for the full subcategory of concrete objects. When $p := \text{cod}$ and $N := N_F$, so that $\text{Krip}(p, N) = \text{Krip}_{\mathcal{M}, F}$, we write $\text{Conc}_{\mathcal{M}, F}$ for $\text{Conc}(p, N)$.

Notation 6.1. We reserve U for forgetful functors into the base category (either $\mathbb{C}$ or $\mathcal{M}$) and indicate the domain by a superscript, so that e.g. $U^{\text{Conc}(p, N)} : \text{Conc}(p, N) \rightarrow \mathbb{C}$.

Example 6.1 (cf. [O’Hearn and Riecke 1995]). Let $F : \mathbb{A} \rightarrow \mathbf{Fin}$. A Kripke relation $(\underline{X}, R) \in \text{Krip}_{\mathbf{Fin}, F}$ (Def. 2.2) is concrete if and only if $\Delta(\Gamma) := \{\lambda\gamma \in F\Gamma . x \mid x \in \underline{X}\} \subseteq R(\Gamma)$ for every $\Gamma \in \mathbb{A}$.

Every Kripke relation has a *concrete completion*: the inclusion $j : \text{Conc}_{\mathcal{M}, F} \hookrightarrow \text{Krip}_{\mathcal{M}, F}$ has a left adjoint defined by $K(\underline{X}, R) = (\underline{X}, KR)$, where $(KR)(\Gamma) := R(\Gamma) \cup \Delta(\Gamma)$ (cf. also (2)). This is an instance of a general construction.

Definition 6.2. Let $(\underline{X}, R) \in \text{Krip}(p, N)$. The set $\text{CE}(\underline{X}, R)$ of *concrete extensions* consists of those $S \in \mathbb{E}$ such that $p(S) = \underline{X}$, $(\underline{X}, R) \leq (\underline{X}, S)$ and $(\underline{X}, S)$ is concrete. The *concrete completion* $(\underline{X}, \tilde{R})$ of $(\underline{X}, R)$ is the fibred product of all the concrete extensions: $(\underline{X}, \tilde{R}) := \bigwedge_{S \in \text{CE}(\underline{X}, R)} (\underline{X}, S)$.

The concrete completion operation extends to a functor K which is left adjoint to j , thereby exhibiting $\text{Conc}(p, N)$ as a reflective subcategory of $\text{Krip}(p, N)$. For $\text{Conc}(p, N)$ to inherit a monad from $\text{Krip}(p, N)$, we further ask for $\text{Conc}(p, N)$ to be a *coreflective* subcategory. This amounts to adding a *restriction* operation, taking an object of $\text{Krip}(p, N)$ to the largest concrete subobject.

Definition 6.3. $\text{Krip}(p, N)$ admits a *hull functor* if the inclusion $j : \text{Conc}(p, N) \hookrightarrow \text{Krip}(p, N)$ has a right adjoint H . We denote the unit and counit of the adjunction $j \dashv H$ by e and c , respectively, and write ϖ for the canonical map witnessing that H preserves products.

Example 6.2. In Ex. 6.1, the hull functor and its counit are given by the inclusion $H\underline{X} \hookrightarrow \underline{X}$:

$$H\underline{X} := \{x \in \underline{X} \mid \lambda\gamma . x \in R(\Gamma) \text{ for every } \Gamma \in \mathbb{A}\}, \quad h \in (HR)\Gamma \iff (F\Gamma \xrightarrow{h} H\underline{X} \hookrightarrow \underline{X}) \in R(\Gamma).$$

In the preceding example, and in those we consider in Sec. 11, the action of the hull functor on relations is determined by the counit c . We axiomatise this with the next definition.

Definition 6.4. A hull functor H is *tractable* if the diagram below is a cartesian lifting for all $(\underline{X}, R)$ in $\text{Krip}(p, N)$. We denote H ’s action by $(\underline{X}, R) \mapsto (H\underline{X}, HR)$ and the counit’s components by $(c, \hat{c})$.

$$\begin{array}{ccc} jHR & \xrightarrow{\hat{c}(\underline{X}, R)} & R \\ & & \downarrow p \\ NjH\underline{X} & \xrightarrow{Nc(\underline{X}, R)} & N\underline{X} \\ & & \downarrow \varpi \\ & & B \end{array}$$

Example 6.3. Let $F : \mathbb{A} \rightarrow \mathcal{M}$. The hull $H : \text{Krip}_{\mathcal{M}, F} \rightarrow \text{Conc}_{\mathcal{M}, F}$ is tractable iff $h \in (HR)(\Gamma) \iff c_X \circ h \in R(\Gamma)$ for all $(\underline{X}, R) \in \text{Krip}_{\mathcal{M}, F}$. Hull functors over $\mathbf{Fin}$ are always tractable (recall Ex. 6.2).

Remark 6.1. In the examples considered in this paper the (tractable) hull functor is always defined similarly to Ex. 6.2. This is because our models $\mathcal{M}$ are categories of sets-with-structure such that, if $X \in \mathcal{M}$ has carrier set $|X| \in \mathbf{Set}$, and $S \subseteq |X|$, then there exists a mono $\tilde{S} \hookrightarrow X$ in $\mathcal{M}$ such that $\tilde{S}$ has carrier S . In future work we shall seek a more widely-applicable condition (see Sec. 1.4).

When $\text{Krip}(p, N)$ admits a hull functor the subcategory $\text{Conc}(p, N)$ becomes a bi-CCC by the theory of (co)reflective subcategories (e.g. Adamek et al. [2009]; Borceux [1994]). Products and coproducts are inherited from $\text{Krip}(p, N)$, and exponentials are given as follows:

$$(X \Rightarrow_{\text{Conc}(p, N)} Y) := H(jX \Rightarrow_{\text{Krip}(p, N)} jY), \quad \text{eval}^{\text{Conc}(p, N)} := \text{eval}^{\text{Krip}(p, N)} \circ (c_{jX \Rightarrow jY} \times jX)$$

Moreover, a strong monad $(\hat{T}, \hat{\mu}, \hat{\eta}, \hat{\text{st}})$ on $\text{Krip}(p, N)$ defines a monad $(W, \mu^W, \eta^W, \text{st}^W)$ on $\text{Conc}(p, N)$:

$$W := H\hat{T}j, \quad \mu^W := H\hat{T}\hat{\mu}j \circ H\hat{T}c\hat{T}j, \quad \eta^W := H\hat{\eta}j \circ e, \quad \text{st}^W := H\hat{\text{st}} \circ \varpi \circ (e \times H\hat{T}j).$$

Lemma 6.1. *If $\text{Krip}(p, N)$ admits a hull functor, then (j, c) is a morphism of strong monads $W \Rightarrow \hat{T}$:*

$$c_{\hat{T}j} \circ \mu^W = \mu^{\hat{T}} \circ \hat{T}c_{\hat{T}j} \circ c_{\hat{T}jW}, \quad c_{\hat{T}j} \circ \eta^W = \eta^{\hat{T}}, \quad c_{\hat{T}j} \circ \text{st}^W = \text{st}^{\hat{T}} \circ (\text{id} \times c_{\hat{T}j}).$$

Hence, if c is component-wise monic, the monadic structure of W is determined by that of $\hat{T}$.

The next result summarises this section.

Proposition 6.1. *For any bi-CCC $\mathbb{C}$, fibration for logical relations $p : \mathbb{E} \rightarrow \mathbb{B}$ and product-preserving functor $N : \mathbb{C} \rightarrow \mathbb{B}$ such that $\text{Krip}(p, N)$ admits a hull functor, $\text{Conc}(p, N)$ is a bi-CCC and j strictly preserves products and coproducts. Moreover, if $\hat{T}$ is a strong monad on $\text{Krip}(p, N)$ then $H\hat{T}j$ is the underlying functor of a strong monad W on $\text{Conc}(p, N)$, yielding the diagram below:*

$$W := H\hat{T}j \begin{array}{c} \curvearrowright \\ \curvearrowleft \end{array} \text{Conc}(p, N) \begin{array}{c} \xleftarrow{K} \\ \xrightarrow{H} \end{array} \text{Krip}(p, N) \begin{array}{c} \xleftarrow{\hat{T}} \\ \xrightarrow{\hat{T}} \end{array} \quad (5)$$

We emphasise that the inclusion j does *not* preserve exponentials. Instead, the exponential in $\text{Conc}(p, N)$ is related to that in $\text{Krip}(p, N)$ by the counit c . In the presence of slightly more structure—e.g. when p is the subobject fibration—the global elements of the carrier of the exponential in $\text{Conc}(p, N)$ may be identified with morphisms in $\mathcal{M}$ preserving the relevant Kripke relations.

Lemma 6.2. *Suppose that, in the situation of Prop. 6.1, it is moreover the case that for every $A \in \mathbb{B}$ the (partially-ordered) fibre $\mathbb{E}_A$ over A has a bottom element M_A . Then $U^{\text{Krip}(p, N)}$ has a left adjoint $L : \mathbb{C} \rightarrow \text{Krip}(p, N) : A \mapsto (A, M_A)$. The composite $K \circ L$ preserves the terminal object, and hence determines a natural isomorphism $\mathcal{M}(1, U^{\text{Conc}(p, N)}(X \Rightarrow_{\text{Conc}(p, N)} Y)) \cong \text{Krip}(p, N)(jX, jY)$.*

7 ABSTRACT OHR CONSTRUCTIONS FOR λ_c

We can now execute the strategy outlined in Sec. 2. We start with the *abstract* construction, in which we assume the properties we need on the indexing set; in Sec. 8 we show how to choose this set concretely. Fix a $\lambda_c(\mathcal{S})$ -model $(\mathcal{M}, T, s)$ and a small set $\mathbb{I}$ such that for each $i \in \mathbb{I}$ one has:

$$\text{A small category } \mathbb{A}_i, \text{ a functor } F_i : \mathbb{A}_i \rightarrow \mathcal{M}, \text{ and a monad lifting } \hat{T}_i \text{ of } T \text{ to } \text{Krip}_{\mathcal{M}, F_i}. \quad (6)$$

For each $i \in \mathbb{I}$ one has $\text{Krip}_{\mathcal{M}, F_i} := \text{Krip}(\text{cod}, N_{F_i})$ and its subcategory of concrete objects $\text{Conc}_{\mathcal{M}, F_i} := \text{Conc}(\text{cod}, N_{F_i})$. Similarly, taking the fibration for logical relations $\prod_{i \in \mathbb{I}} \text{cod}$ one obtains a category $\mathcal{K} := \text{Krip}(\prod_{i \in \mathbb{I}} \text{cod}, \langle N_{F_i} \rangle_{i \in \mathbb{I}})$ with objects denoted $X := (\underline{X}, \bar{X})$ as in Sec. 2, and a subcategory $\mathcal{C} := \text{Conc}(\prod_{i \in \mathbb{I}} \text{cod}, \langle N_{F_i} \rangle_{i \in \mathbb{I}})$. The structure in $\mathcal{K}$ and $\mathcal{C}$ is determined component-wise.

Lemma 7.1. (1) *Take $(\underline{Y}, \bar{Y}) \in \mathcal{K}$, $f : \underline{X} \rightarrow \underline{Y}$ in $\mathcal{M}$, and cartesian liftings in $\mathcal{K}$ and $\text{Krip}_{\mathcal{M}, F_i}$:*

$$\begin{array}{ccccc} \prod_{i \in \mathbb{I}} \text{Sub}(\widehat{\mathbb{A}}_i) & f^*(\bar{Y}) \dashrightarrow & \bar{Y} & f^*(\bar{Y}(i)) \dashrightarrow & \bar{Y}(i) & \text{Sub}(\widehat{\mathbb{A}}_i) \\ \Pi_i \text{cod} \downarrow & & & & & \downarrow \text{cod} \\ \prod_{i \in \mathbb{I}} \widehat{\mathbb{A}}_i & \langle N_{F_i} \rangle_i(\underline{X}) \xrightarrow{\langle N_{F_i} \rangle_i(f)} & \langle N_{F_i} \rangle_i(\underline{Y}) & N_{F_i}(\underline{X}) \xrightarrow{N_{F_i}(f)} & N_{F_i}(\underline{Y}) & \widehat{\mathbb{A}}_i \end{array}$$

Then $f^(\bar{Y})$ is determined component-wise, in the sense that $(f^*(\bar{Y}))(i) = f^*(\bar{Y}(i))$ for all $i \in \mathbb{I}$.*

(2) *Setting $\hat{T}(\underline{X}, \bar{X}) := (T\underline{X}, \hat{T}\bar{X})$, where $\hat{T}\bar{X}(i) = \hat{T}_i(\bar{X}(i))$ for each $i \in \mathbb{I}$, defines a lifting of T to $\mathcal{K}$.*

(3) *$(\underline{X}, \bar{X}) \in \mathcal{C}$ if and only if $(\underline{X}, \bar{X}(i)) \in \text{Conc}_{\mathcal{M}, F_i}$ for every $i \in \mathbb{I}$.*

(4) *$f : \underline{X} \rightarrow \underline{Y}$ in $\mathcal{C}$ if and only if $f : (\underline{X}, \bar{X}(i)) \rightarrow (\underline{Y}, \bar{Y}(i))$ in $\text{Conc}_{\mathcal{M}, F_i}$ for every $i \in \mathbb{I}$.*

For our abstract construction we want to assume enough structure so that Prop. 6.1 holds. This is captured by the following; assumption (3) axiomatises the situation of Ex. 6.2.

Assumption 7.1. We assume the following: (1) an interpretation $\hat{s}$ of base types, operations and primitives with $U^C \circ \hat{s} = s$; (2) an index $i_0 \in \mathbb{I}$ with $\mathbb{A}_{i_0} = \text{Cons}_s$ and $F_{i_0} := U^C \circ \hat{s}[-]$. Moreover, we assume: (3) $\mathcal{K}$ admits a tractable hull functor H with counit $c : jH \Rightarrow \text{id}$ component-wise monic.

With these assumptions, Prop. 6.1 entails that C acquires a bi-CCC structure and a strong monad W with underlying functor $H\hat{T}j$. We therefore recover the situation in Figure 1b.

Definition 7.1. We call $(C, W, \hat{s})$ the *abstract OHR model* on $(\mathcal{M}, T, s)$.

7.1 $\lambda_c(\mathcal{S})$ -Logical Relations Over C

Having constructed $(C, W, \hat{s})$ we turn to considering logical relations over this model and relating them to logical relations over $(\mathcal{M}, T, s)$. To this end we consider the following two diagrams, in which $\hat{W}$ is any lifting of W ; recall that $\text{Krip}(\text{cod}, N_{U \circ \hat{s}}[-])$ is exactly $\text{Krip}_{\mathcal{M}, F_{i_0}}$ by Assump. 7.1(2).

$$\begin{array}{ccc} \hat{W} \hookrightarrow \text{Krip}_{C, \hat{s}[-]} & \longrightarrow & \text{Sub}(\widehat{\text{Cons}}) \\ \downarrow U & \lrcorner & \downarrow \text{cod} \\ W \hookrightarrow C & \xrightarrow{N_{\hat{s}[-]}} & \widehat{\text{Cons}} \end{array} \quad \begin{array}{ccc} \hat{T}_{i_0} \hookrightarrow \text{Krip}_{\mathcal{M}, F_{i_0}} & \longrightarrow & \text{Sub}(\widehat{\text{Cons}}) \\ \downarrow \lrcorner & & \downarrow \text{cod} \\ T \hookrightarrow \mathcal{M} & \xrightarrow{N_{U \circ \hat{s}}[-]} & \widehat{\text{Cons}} \end{array} \quad (7)$$

As observed in Sec. 2, for any $(\underline{X}, \mathcal{R}) \in \text{Krip}_{C, \hat{s}[-]}$ the faithfulness of U in (7) yields a chain of inclusions $\mathcal{R} \hookrightarrow C(\hat{s}[-], X) \hookrightarrow \mathcal{M}(U\hat{s}[-], \underline{X})$, so $(\underline{X}, \mathcal{R})$ becomes an object in $\text{Krip}_{\mathcal{M}, F_{i_0}}$.

Remark 7.1. The bi-CCC structure of $\text{Krip}_{C, \hat{s}[-]}$ is given as in Sec. 4.1, except one must take care to use the bi-CCC structure of C . For example, the exponential $(X, \mathcal{R}) \Rightarrow_{\text{Krip}_{C, \hat{s}[-]}} (Y, \mathcal{S})$ is $(H(X \Rightarrow Y), \overline{H(X \supset Y)}, \mathcal{R} \ni \mathcal{S})$, where $h \in (\mathcal{R} \ni \mathcal{S})\Gamma \iff \text{eval} \circ \langle c_{jX \Rightarrow jY} \circ h \circ \hat{s}[\rho], u \rangle \in \mathcal{S}(\Delta)$ for any $\rho : \Gamma \rightarrow \Delta$ and $u \in \mathcal{R}(\Gamma)$; equivalently, $c_{jX \Rightarrow jY} \circ h \in (\mathcal{R} \supset \mathcal{S})(\Gamma)$.

Together with the fact the forgetful functors $\text{Krip}_{C, \hat{s}[-]} \rightarrow C$ and $C \rightarrow \mathcal{M}$ both strictly preserve products, Remark 7.1 entails that—so long as $\hat{W}$ interacts well with $\hat{T}_{i_0}$ —then $\lambda_c(\mathcal{S})$ -logical relations over $\hat{W}$ are ‘tracked’ by the cartesian closed structure of $\text{Krip}_{\mathcal{M}, F_{i_0}}$. The proof is by induction.

Lemma 7.2. Let $\mathcal{R} = \{\mathcal{R}_\sigma\}_{\sigma \in \text{Ty}}$ be a $\lambda_c(\mathcal{S})$ -logical relation over $\hat{W}$ such that: (1) $h \in \mathcal{R}_\sigma(\Gamma)$ if and only if $c_{\hat{T}j\hat{s}[\sigma]} \circ h \in (\hat{T}_{i_0}\mathcal{R}_\sigma^{\text{val}})(\Gamma)$; and (2) $\overline{\hat{s}[\sigma]}(i_0) = \mathcal{R}_\beta^{\text{val}}$ for every $\beta \in \text{Ty}$. Then $\overline{\hat{s}[\sigma]}(i_0) = \mathcal{R}_\sigma^{\text{val}}$ and $(\overline{W\hat{s}[\sigma]})(i_0) = \mathcal{R}_\sigma$ for every type $\sigma \in \text{Ty}$.

We now want to find $\hat{W}$ and $\hat{T}_{i_0}$ so that condition (1) in this lemma holds automatically. Recalling Prop. 5.1, it is natural to use the $\top\top$ -lifting arising from the given logical relation, namely $\hat{W} = W^{\top\top[\mathcal{R}]}$. It remains to identify a suitable choice of $\hat{T}_{i_0}$. To this end, note that every $(W\hat{s}[\sigma], \mathcal{R}_\sigma)$ gives rise to a $\top\top$ -lifting parameter $(T\hat{s}[\sigma], \langle \mathcal{R}_\sigma \rangle)$ in $\text{Krip}_{\mathcal{M}, F_{i_0}}$ as follows:

$$\langle \mathcal{R}_\sigma \rangle(\Gamma) := \{c_{\hat{T}j\hat{s}[\sigma]} \circ h \mid h \in \mathcal{R}_\sigma(\Gamma)\} \subseteq \mathcal{M}(\underline{\hat{s}[\Gamma]}, T\hat{s}[\sigma])$$

Lemma 7.3. Let $\mathcal{R} = \{\mathcal{R}_\sigma\}_{\sigma \in \text{Ty}}$ be a hungry $\lambda_c(\mathcal{S})$ -logical relation over $\hat{W} := W^{\top\top[\mathcal{R}]}$. Then condition (1) of Lemma 7.2 holds: $h \in \mathcal{R}_\sigma(\Gamma) \iff c_{\hat{T}j\hat{s}[\sigma]} \circ h \in (T^{\top\top[\langle \mathcal{R} \rangle]}\mathcal{R}_\sigma^{\text{val}})(\Gamma)$.

Combining Lemmas 7.2 and 7.3 yields the following, which makes precise the sense in which morphisms in the abstract OHR model preserve every compatible logical relation over C .

Proposition 7.1. Let $\mathcal{R} = \{\mathcal{R}_\sigma\}_{\sigma \in \text{Ty}}$ be a hungry $\lambda_c(\mathcal{S})$ -logical relation over $\hat{W} := W^{\top\top[\mathcal{R}]}$ and suppose $i_0 \in \mathbb{I}$ is such that $\hat{T}_{i_0} = T^{\top\top[\langle \mathcal{R} \rangle]}$ and $\overline{\hat{s}[\sigma]}(i_0)(\beta) = \mathcal{R}_\beta^{\text{val}}$ for every $\beta \in \mathbf{B}$. Then $\overline{\hat{s}[\sigma]}(i_0) = \mathcal{R}_\sigma^{\text{val}}$ and $(\overline{W\hat{s}[\sigma]})(i_0) = \mathcal{R}_\sigma$ for every type $\sigma \in \text{Ty}$, so every $f : \hat{s}[\Gamma] \rightarrow W\hat{s}[\sigma]$ in C satisfies $\mathcal{R}$.

8 A FULLY ABSTRACT MODEL FOR λ_c

Choosing the Indexing Set $\mathbb{I}$. We now follow the strategy sketched in Sec. 2.2. To instantiate the abstract OHR construction (Def. 7.1) we need to choose $\mathbb{I}$, $i_0 \in \mathbb{I}$, and $\hat{s}$ so that Assump. 7.1 and the hypotheses of Prop. 7.1 hold. Thus, we replace Assump. 7.1 with the weaker assumptions below.

Assumption 8.1. Let $(\mathcal{M}, T, s)$ be a $\lambda_c^+(\mathcal{S})$ -model such that $\mathcal{M}$ is *small* and, for any small set $\mathbb{J}$ and $\mathbb{J}$ -indexed set of functors $\mathbb{A}_j \rightarrow \mathcal{M}$, the category $\text{Krip}(\prod_{j \in \mathbb{J}} \text{cod}, \langle N_{F_j} \rangle_{j \in \mathbb{J}})$ admits a tractable hull functor with counit c component-wise monic.

Remark 8.1. The size restriction is not onerous. Although many models of interest are large (e.g. **Set**, ωCpo , presheaf categories), one generally works within a small subcategory. For example, one may replace **Set** with the subcategory Set_κ of *hereditarily- κ sets*, for κ some infinite cardinal.

We construct $\mathbb{I}$ in stages, ranging over the data required to construct a λ_c^+ -model structure on each category of Kripke relations $\mathcal{K}_i$. By Prop. 5.1 this amounts to ranging over a collection of $\lambda_c(\mathcal{S})$ -logical relations. First we give the data necessary to construct $\mathcal{K}_i$:

- Fix a set $\mathbb{S}$ of small categories such that $\text{Con}_\mathbb{S}^{\text{op}} \in \mathbb{S}$ (the singleton $\{\text{Con}_\mathbb{S}^{\text{op}}\}$ suffices).
- For each $\mathbb{A} \in \mathbb{S}$, let $\text{Fun}(\mathbb{A}) := [\mathbb{A}, \mathcal{M}]$ be the set of functors $\mathbb{A} \rightarrow \mathcal{M}$; this is small since $\mathcal{M}$ is.

Since the nerve functor $N_F : \mathcal{M} \rightarrow \hat{\mathcal{A}}$ preserves limits for any $F : \mathbb{A} \rightarrow \mathcal{M}$, for each $\mathbb{A} \in \mathbb{S}$ and $F \in \text{Fun}(\mathbb{A})$ one obtains a bi-CCC $\text{Krip}_{\mathcal{M}, F}$. We then take

- For each $\mathbb{A} \in \mathbb{S}$ and $F \in \text{Fun}(\mathbb{A})$, take $\text{Lift}(\mathbb{A}, F)$ to be the set of liftings of T to $\text{Krip}_{\mathcal{M}, F}$. This is small because $\mathcal{M}$ is small and $\text{Krip}_{\mathcal{M}, F} \rightarrow \mathcal{M}$ has small fibres.

Clearly $U \circ \hat{s}[-] \in \text{Fun}(\text{Con}_\mathbb{S}^{\text{op}})$, so Assump. 7.1(2) will hold, and $T^{\top\top}[\langle \text{DEF} \rangle] \in \text{Lift}(\text{Con}_\mathbb{S}^{\text{op}}, U\hat{s}[-])$.

It remains to construct the interpretation $\hat{s}$. By Lemma 7.1 it suffices to work component-wise, so fix $\mathbb{A} \in \mathbb{S}$, $F \in \text{Fun}(\mathbb{A})$ and $\hat{T} \in \text{Lift}(\mathbb{A}, F)$. We need each base type to be interpreted by a concrete object, so we range over interpretations of base types $r : \mathbf{B} \rightarrow \text{Conc}_{\mathcal{M}, F}$ lying over the interpretation s . However, to interpret primitives with thunks we want to use the exponentials of $\text{Krip}_{\mathcal{M}, F}$, which lie over those in $\mathcal{M}$, rather than those of $\text{Conc}_{\mathcal{M}, F}$. Thus, we use the composite $j \circ r : \mathbf{B} \rightarrow \text{Krip}_{\mathcal{M}, F}$. This extends canonically to an interpretation $(j \circ r)[-]$ of all types, so we take

- $\text{Interp}(\mathbb{A}, F, \hat{T})$ is the set of all maps $r : \mathbf{B} \rightarrow \text{Conc}_{\mathcal{M}, F}$ such that: (1) $U^{\text{Conc}_{\mathcal{M}, F}} \circ r = s$; and (2) $s[\text{op}] : jr[\alpha] \rightarrow \hat{T}(jr[\kappa])$ and $s[\xi] : 1 \rightarrow (jr)[\kappa]$ for each $\text{op} : \alpha \rightsquigarrow \kappa$ and $\xi : \kappa$ in $\mathcal{S}$.

This set is small because the fibration $\text{Krip}_{\mathcal{M}, F} \rightarrow \mathcal{M}$ has small fibres, so for each $\beta \in \mathbf{B}$ there's a small set of choices of object lying over $s[\beta]$. Putting everything together, we define

$$\mathbb{I} := \{(\mathbb{A}, F, \hat{T}, r) \mid \mathbb{A} \in \mathbb{S}, F \in \text{Fun}(\mathbb{A}), \hat{T} \in \text{Lift}(\mathbb{A}, F), r \in \text{Interp}(\mathbb{A}, F, \hat{T})\} \quad (8)$$

We now detail the OHR construction. Following O'Hearn and Riecke [1995], set $\hat{s}(\beta) = (s(\beta), \overline{\hat{s}(\beta)})$ where $\overline{\hat{s}(\beta)}(\mathbb{A}, F, \hat{T}, r)$ is the second projection of $r(\beta)$; this is a mapping $\mathbf{B} \rightarrow \mathcal{C}$ by Lemma 7.1(3). To interpret operations and primitives we use the following lemma, which shows that these lift from $\mathcal{M}$ to $\mathcal{K}$; we can then use the adjunction $j \dashv H$ to construct the required maps in $\mathcal{C}$. Thunks require particular care because if $\sigma_i = (1 \rightarrow \gamma_i)$ then $\hat{s}[\sigma_i] = H(1 \Rightarrow jW\hat{s}[\gamma_i])$. We therefore define an object $\ulcorner \hat{s}[\sigma_i] \urcorner$ and map $m_{\sigma_i} : j\hat{s}[\sigma_i] \rightarrow \ulcorner \hat{s}[\sigma_i] \urcorner$ by setting $\ulcorner \hat{s}[\sigma_i] \urcorner := \hat{s}[\gamma_i]$ and $m_{\sigma_i} := \text{id}$ if $\sigma_i = \gamma_i$, and $\ulcorner \hat{s}[\sigma_i] \urcorner := 1 \Rightarrow \hat{T}j\hat{s}[\gamma_i]$ and $m_{\sigma_i} := (\text{id} \Rightarrow c_{\hat{T}j\hat{s}[\gamma_i]}) \circ c_{1 \Rightarrow W\hat{s}[\gamma_i]}$ if σ_i is a thunk $(1 \rightarrow \gamma_i)$.

Lemma 8.1. For $\mathbb{I}$ and $\hat{s} : \mathbf{B} \rightarrow \mathcal{C}$ as defined above, and $\mathcal{K}$ and $\mathcal{C}$ constructed as in Figure 1b:

- (1) if $\text{op} : \alpha \rightsquigarrow \kappa$, then $s[\text{op}] : j\hat{s}[\alpha] \rightarrow \hat{T}\hat{s}[\kappa]$ in $\mathcal{K}$; and (2) if $\xi : \gamma$, then $s[\xi] : 1 \rightarrow \hat{s}[\gamma]$ in $\mathcal{C}$; and
- (3) if $\xi : \sigma_1 * \dots * \sigma_n \rightarrow \gamma$, then $s[\xi] : 1 \rightarrow (j\ulcorner \hat{s}[\sigma_1] \urcorner \times \dots \times j\ulcorner \hat{s}[\sigma_n] \urcorner) \Rightarrow \hat{T}\hat{s}[\gamma]$ in $\mathcal{K}$.

For the interpretation of primitives note that for any category $\text{Krip}(p, N)$ with hull functor H there is a bijective correspondence between maps $1 \rightarrow H(jX \Rightarrow jH\hat{T}jY)$ and maps $1 \times jX \rightarrow \hat{T}jY$.

Definition 8.1. The *OHR model* $(\text{OHR}(\mathcal{M}), W, \hat{s})$ over $(\mathcal{M}, T, s)$ is the $\lambda_c(S)$ -model obtained by instantiating Def. 7.1 with: (1) indexing set $\mathbb{I}$ as in (8), and $\mathbb{A}_i, F_i$ and $\hat{T}_i$ given by the projections; and (2) interpretation $\hat{s}$ with $\hat{s}(\beta) := s(\beta)$ and $\hat{s}(\beta)(\mathbb{A}, F, \hat{T}, r)$ the second projection of $r(\beta)$; and (3) operations $\text{op} : \alpha \rightsquigarrow \kappa$ interpreted by $\hat{s}[\text{op}] := H(s[\text{op}]) \circ e_{s[\alpha]}$, primitives $(\xi : \gamma)$ by $\hat{s}[\xi] := s[\xi]$ and primitives $(\xi : \sigma_1 * \dots * \sigma_n \rightarrow \gamma)$ by setting $\hat{s}[\xi]$ to be the arrow corresponding to $\text{eval} \circ (s[\xi] \times \prod_{i=1}^n m_{\sigma_i})$ across the bijective correspondence above.

The next result makes precise the idea that $\text{OHR}(\mathcal{M})$ -morphisms preserve every compatible logical relation over $\text{OHR}(\mathcal{M})$. For the proof, set $i_0 := (\text{Con}_S^{\text{op}}, U \circ \hat{s}[-], T^{\text{TT}}[\langle \mathcal{R} \rangle], r)$ with $r : \beta \mapsto (s[\beta], R_\beta^{\text{val}})$, show that $i_0 \in \mathbb{I}$, i.e. that $r \in \text{Interp}(\text{Con}_S^{\text{op}}, U \circ \hat{s}[-], T^{\text{TT}}[\langle \mathcal{R} \rangle])$, and apply Prop. 7.1.

Proposition 8.1. Let $\mathcal{R} = \{\mathcal{R}_\sigma\}_{\sigma \in \text{Ty}}$ be a hungry, λ -compatible $\lambda_c(S)$ -logical relation over $\hat{W} := W^{\text{TT}}[\mathcal{R}]$ such that every global element $g : 1 \rightarrow s[\beta]$ in $\mathcal{M}$ is in $\mathcal{R}_\beta^{\text{val}}(\diamond)$. Then any morphism $f : \hat{s}[\Gamma] \rightarrow W\hat{s}[\sigma]$ in $\text{OHR}(\mathcal{M})$ satisfies $\mathcal{R}$, and hence—since $\mathcal{R}$ is hungry—is in $\mathcal{R}_\sigma(\Gamma)$.

Together with Prop. 5.1, the preceding entails $\text{OHR}(\mathcal{M})$ is “saturated”: if there exist enough global elements, no $\lambda_c(S)$ -morphism can cut out any morphisms. This should be contrasted with Ex. 1.2.

Corollary 8.1. If $F : (\mathcal{N}, S, t) \rightarrow (\text{OHR}(\mathcal{M}), W, \hat{s})$ is any strict $\lambda_c(S)$ -morphism such that the induced composite $(U^{\text{OHR}(\mathcal{M})} \circ F)_{1,t}[\beta] : \mathcal{N}(1, t[\beta]) \rightarrow \mathcal{M}(1, s[\beta])$ is onto, then every $f : \hat{s}[\Gamma] \rightarrow W\hat{s}[\sigma]$ in $\text{OHR}(\mathcal{M})$ is in the image of F .

Full Completeness of the OHR Construction. Since DEF is always $\lambda_c(S)$ -logical over $T^{\text{TT}}[\text{DEF}]$ (Thm. 5.1), the final obstacle to full completeness is the concreteness condition on global elements. For this we relate the induced interpretation $\hat{s}$ to s . Write $\mathcal{D}_\sigma(\Gamma) := \{f \mid f \text{ is definable in } (\mathcal{M}, T, s)\}$ for any context Γ and $\sigma \in \text{Ty}$, and call a context Γ *ground* if $\sigma \in \mathbf{G}$ whenever $(x : \sigma) \in \Gamma$. Since the forgetful functor $U^{\text{OHR}(\mathcal{M})}$ strictly preserves products, $\hat{s}[\Gamma] = s[\Gamma]$ for any ground context Γ .

Lemma 8.2. For every ground context Γ and ground type γ , $U^{\text{OHR}(\mathcal{M})}(c_{\hat{T}j\hat{s}[\gamma]} \circ \hat{s}[\Gamma \vdash M : \gamma]) = s[\Gamma \vdash M : \gamma]$. Hence DEF and $\mathcal{D}$ coincide on closed terms of ground type: $\text{DEF}_Y^{\text{val}}(\diamond) = \mathcal{D}_Y^{\text{val}}(\diamond)$.

Remark 8.2. It follows that our semantic definition of contextual equivalence is consistent between $(\mathcal{M}, T, s)$ and $(\text{OHR}(\mathcal{M}), W, \hat{s})$: we have $M \simeq_{\text{ctx}} M'$ in $\mathcal{M}$ if and only $M \simeq_{\text{ctx}} M'$ in $\text{OHR}(\mathcal{M})$.

By Lemma 8.2, if a global element $g : 1 \rightarrow s[\beta]$ in $\mathcal{M}$ is such that $\eta_{s[\beta]} \circ g$ is definable, then this composite is also definable in $\text{OHR}(\mathcal{M})$. Hence from Prop. 8.1 we obtain full completeness.

Theorem 8.1. Let $(\mathcal{M}, T, s)$ be a small, well-pointed $\lambda_c(S)$ -model such that for every $\beta \in \mathbf{B}$ and global element $g : 1 \rightarrow s[\beta]$ the composite $\eta_{s[\beta]} \circ g$ is definable. Then, if $\text{Krip}(\prod_i \text{cod}, \langle F \rangle_{i \in \mathbb{I}})$ admits a tractable hull functor and the counit of the adjunction $j \dashv H$ is component-wise monic, the induced model $(\text{OHR}(\mathcal{M}), W, \hat{s})$ is well-pointed and fully complete, hence fully abstract.

Example 8.1. The theorem applies to any model on **Set** in which the signature has a primitive $\underline{b} : \beta$ for every $b \in s(\beta)$ and $\beta \in \mathbf{B}$ (for the size restriction take Set_κ for a large enough κ , cf. Remark 8.1), such as when one has a base type nat and primitives $\underline{n} : \text{nat}$ for every $n \in \mathbb{N}$.

9 EXAMPLE: A FULLY ABSTRACT MODEL FOR IMMUTABLE STATE

The structure of the OHR model is closely related to that of the original model, and one can use this to compute in the OHR model. We highlight these properties by returning to Ex. 1.1: let us denote

the signature considered there by S_{RO} and write $(\mathbf{Fin}, R, s)$ for the associated semantic model. We shall show why the counterexample κ , which shows that $(\mathbf{Fin}, R, s)$ is not fully abstract, does not give rise to a counterexample in $(\mathbf{OHR}(\mathbf{Fin}), W, \hat{s})$; along the way we shall see a particular example of how Prop. 8.1 ensures full completeness of the OHR model.

We begin by making explicit the structure of the OHR model for a $\lambda_c(S)$ -model $(\mathcal{M}, T, s)$ with $\mathcal{M} \subseteq \mathbf{Set}$. Recalling Ex. 6.2, the counit c is an inclusion on carrier sets: for any $X = (\underline{X}, \bar{X}) \in \mathcal{K}$ one has $c_X : jHX \hookrightarrow \underline{X}$. All the λ_c -model structure of $\mathbf{OHR}(\mathcal{M})$ is then determined by c and the corresponding structure on $\mathcal{M}$. For example, products in $\mathbf{OHR}(\mathcal{M})$ coincide with products in $\mathcal{M}$ and $\text{eval}^{\mathbf{OHR}(\mathcal{M})}$ is the following composite in $\mathcal{M}$, so that $\text{eval}^{\mathbf{OHR}(\mathcal{M})}(f, x) = f(x)$:

$$\underline{H(jX \Rightarrow jY) \times X} \hookrightarrow (\underline{X \Rightarrow Y}) \times \underline{X} \xrightarrow{\text{eval}} \underline{Y} \quad (9)$$

By Lemma 6.1, similar remarks apply to the monadic structure, e.g. $\eta_X^W(x) = \eta_X^T(x)$ for all $x \in \underline{X}$.

We now return to the particular case of Ex. 1.1. Since $s[\text{tt}]$ and $s[\text{ff}]$ name the two elements of $s[\text{bool}]$, the model $(\mathbf{Fin}, R, s)$ satisfies the conditions of Thm. 8.1. Hence the OHR model $(\mathbf{OHR}(\mathbf{Fin}), W, \hat{s})$ on $(\mathbf{Fin}, R, s)$ exists and is fully abstract.

Using the relationship between $\mathbf{OHR}(\mathbf{Fin})$ and $\mathbf{Fin}$ sketched above, one sees the following.

Example 9.1. For any closed $\lambda_c(S_{RO})$ -terms $N, N' : 1 \rightarrow \text{bool}$ and variable $f : (1 \rightarrow \text{bool}) \rightarrow \text{bool}$,

$$c_{\hat{T}j\hat{s}[\text{bool}]} \circ \hat{s}[f : (1 \rightarrow \text{bool}) \rightarrow \text{bool} \vdash \text{or}(\langle f N, f N' \rangle) : \text{bool}]] = \lambda\varphi. \lambda i. \varphi(n(i))(i) \vee \varphi(n'(i))(i)$$

where $n = c_{\hat{T}j\hat{s}[1 \rightarrow \text{bool}]} \circ \hat{s}[N]$ and $n' = c_{\hat{T}j\hat{s}[1 \rightarrow \text{bool}]} \circ \hat{s}[N']$. Since the components of c are injections, this determines $\hat{s}[\text{or}(\langle f N, f N' \rangle)]$. Similar considerations show that, where M and M' are as in (1), then the set maps $\hat{s}[M]$ and $\hat{s}[M']$ have the same action as $s[M]$ and $s[M']$.

We can now show why the counterexample morphism κ of Ex. 1.1 cannot determine an element of $\hat{s}[(1 \rightarrow \text{bool}) \rightarrow \text{bool}]$. The idea, which holds quite generally, is to instantiate a version of Figure 1a with $\mathbf{Fin}$ replaced by $\mathbf{OHR}(\mathbf{Fin})$, then adapt the argument from Ex. 1.2 to show κ does not restrict to a map in $\mathbf{OHR}(\mathbf{Fin})$. Accordingly, let $\mathbb{L}_{\mathbf{OHR}(\mathbf{Fin})}$ be the category obtained by change-of-base along $\Delta \circ U^{\mathbf{OHR}(\mathcal{M})} : \mathbf{OHR}(\mathcal{M}) \rightarrow \mathbf{Fin} \times \mathbf{Fin}$. Explicitly, $\mathbb{L}_{\mathbf{OHR}(\mathbf{Fin})}$ has objects triples $((\underline{X}, \bar{X}), R_1, R_2)$ where each R_i is a binary relation on the carrier $\underline{X}$; products and exponentials are defined as in $\mathbb{L}_{\mathbf{Fin}}$.

For the monad, the universal property of the fibration induces a lifting $\hat{W}$ of W to $\mathbb{L}_{\mathbf{OHR}(\mathbf{Fin})}$: for $((\underline{X}, \bar{X}), S) \in \mathbb{L}_{\mathbf{OHR}(\mathbf{Fin})}$ one has $(\underline{X}, S) \in \mathbb{L}_{\mathbf{Fin}}$ and $(R\underline{X}, \hat{R}S) \in \mathbb{L}_{\mathbf{Fin}}$, so one defines $\hat{W}S$ using the cartesian lifting in (10), below. Explicitly, $(h, h') \in (\hat{W}R)_i \iff (h i, h' i) \in R_i$. To finish defining a the semantic model, set $\hat{t}(\text{bool}) = (\hat{s}[\text{bool}], \{(0, 0), (1, 1)\}, \{(0, 0), (1, 1)\})$ and interpret primitives and operations using the universal property of the cartesian lifting. Then $(\mathbb{L}_{\mathbf{OHR}(\mathbf{Fin})}, \hat{W}, \hat{t})$ is a $\lambda_c(S_{RO})$ -model and the forgetful functor $\mathbb{L}_{\mathbf{OHR}(\mathbf{Fin})} \rightarrow \mathbf{OHR}(\mathbf{Fin})$ is a $\lambda_c(S_{RO})$ -morphism.

$$\begin{array}{ccc} \hat{W}S & \dashrightarrow & \hat{R}S \\ & & \text{Sub}(\mathbf{Fin}) \times \text{Sub}(\mathbf{Fin}) \\ & & \downarrow \text{cod} \times \text{cod} \\ \underline{NWX} & \xrightarrow{\text{Nc}_{\hat{T}jX}} & \underline{NR(X)} & \text{Fin} \times \text{Fin} \end{array} \quad (10)$$

We can now show the counterexample κ from Ex. 1.1 does not restrict to a counterexample κ' in $\mathbf{OHR}(\mathbf{Fin})$, i.e. that there does not exist κ' in $\mathbf{OHR}(\mathbf{Fin})$ and $i \in 2$ such that

$$\kappa = \left(1 \xrightarrow{\kappa'} \underline{W\hat{s}[(1 \rightarrow \text{bool}) \rightarrow \text{bool}]} \hookrightarrow R((1 \Rightarrow R2) \Rightarrow R2) \xrightarrow{\text{eval}_i} (1 \Rightarrow R2) \Rightarrow R2 \right).$$

If such a κ' existed, it must lift to $\mathbb{L}_{\mathbf{OHR}(\mathbf{Fin})}$ by Cor. 8.1, but, arguing in the same way as Ex. 1.2, one sees this is impossible. Similar reasoning remedies our omission in Ex. 1.1.

Lemma 9.1. *For M, M' as in Ex. 1.1, $M \simeq_{\text{ctx}} M'$ in $(\mathbf{Fin}, R, s)$.*

PROOF. By Remark 8.2 it suffices to show $M \simeq_{\text{ctx}} M'$ in the OHR model. Suppose for a contradiction that $\hat{s}[\![M]\!] \neq \hat{s}[\![M']]\!]$. The counit determines an inclusion $\overline{W\hat{s}[\![(1 \rightarrow \text{bool}) \rightarrow \text{bool}] \rightarrow \text{bool}]\!]} \hookrightarrow R(\hat{s}[\![(1 \rightarrow \text{bool}) \rightarrow \text{bool}]\!] \rightarrow R2)$ so there exists some $i \in 2$ and $\delta \in \hat{s}[\![(1 \rightarrow \text{bool}) \rightarrow \text{bool}]\!]$ such that $\hat{s}[\![M]\!](i)(\delta) \neq \hat{s}[\![M']]\!](i)(\delta)$. Using the interpretations of read , tt , ff and $\neg$ in $\text{OHR}(\mathcal{M})$ one sees that $\hat{s}[\![(1 \rightarrow \text{bool}) \rightarrow \text{bool}]\!]$ has four elements, and so must equal $s[\![(1 \rightarrow \text{bool}) \rightarrow \text{bool}]\!]$. A long but basic check then shows that if $s[\![M]\!](i)(\omega) \neq s[\![M']]\!](i)(\omega)$ then ω does not lift to a morphism in $\mathbb{L}_{\mathbf{Fin}}$ so, by Ex. 9.1, δ cannot lift to a morphism in $\mathbb{L}_{\mathbf{Fin}}$. On the other hand, by concreteness and Cor. 8.1, δ defines a morphism $1 \rightarrow \hat{t}[\![(1 \rightarrow \text{bool}) \rightarrow \text{bool}]\!]$ in $\mathbb{L}_{\text{OHR}(\mathbf{Fin})}$. But then $\delta : (X, R_1, R_2) \rightarrow (Y, S_1, S_2)$ in $\mathbb{L}_{\text{OHR}(\mathbf{Fin})}$ implies $\delta : (X, R_1, R_2) \rightarrow (Y, S_1, S_2)$ in $\mathbb{L}_{\mathbf{Fin}}$, contradicting the preceding. $\square$

10 A FULLY ABSTRACT MODEL FOR λ_c^+

We now fold sum types into the development of Secs. 7 and 8. We saw in Sec. 5 that to characterise definability with sum types one needs to restrict to $\hat{T}$ -closed objects. Accordingly, for our OHR construction we restrict not just to concrete objects, but to those that are both concrete and $\hat{T}$ -closed. Write $k : \text{ConcCl}(p, N, \hat{T}) \hookrightarrow \text{Krip}(p, N)$ for the subcategory of concrete, $\hat{T}$ -closed objects. We induce structure on this category as we did for the concrete and $\hat{T}$ -closed subcategories.

Lemma 10.1 (cf. Lemma 4.4 and Prop. 6.1). *In the situation of Prop. 6.1:*

- (1) *The closure operator $\text{id}^{\top\top}[\hat{T}]$ restricts to $\text{Conc}(p, N)$: if (A, R) is concrete, then so is $\text{id}^{\top\top}[\hat{T}](A, R)$. Hence $\text{ConcCl}(p, N, \hat{T})$ is a replete reflective subcategory of $\text{Conc}(p, N)$.*
- (2) *$\text{ConcCl}(p, N, \hat{T})$ is an exponential ideal (e.g. Johnstone [2002, p. 52]) of $\text{Conc}(p, N)$.*
- (3) *If H is tractable, then $H\hat{T}j$ restricts to a strong monad on $\text{ConcCl}(p, N, \hat{T})$.*

Hence, $\text{ConcCl}(p, N, \hat{T})$ is a bi-CCC with cartesian-closed structure inherited from $\text{Conc}(p, N)$ and coproducts as in Lemma 4.4(2), the forgetful functor $U : \text{ConcCl}(p, N, \hat{T}) \rightarrow \mathbb{A}$ strictly preserves finite products and finite coproducts, and one obtains the diagram below (cf. diagrams (4) and (5)):

$$\begin{array}{ccccc} & & H\hat{T}j & & \\ & & \Downarrow & & \\ H\hat{T}k \circlearrowleft & \text{ConcCl}(p, N, \hat{T}) & \xleftarrow{\perp} & \text{Conc}(p, N) & \xleftrightarrow[\perp]{K} \text{Krip}(p, N) \circlearrowright \hat{t} \end{array} \quad (11)$$

We now refine the abstract OHR construction for $\lambda_c(S)$ (Sec. 7) to incorporate sums. Consider a fixed $\lambda_c^+(S)$ -model $(\mathcal{M}, T, s)$ and an indexing set $\mathbb{I}$ such that for each $i \in \mathbb{I}$ one has chosen data as in (6), and make the assumptions of Assump. 7.1. For each $i \in \mathbb{I}$ one now obtains three categories: $\text{Krip}_{\mathcal{M}, F_i} := \text{Krip}(\text{cod}, N_{F_i})$, $\text{Conc}_{\mathcal{M}, F_i} := \text{Conc}(\text{cod}, N_{F_i})$ and $\text{ConcCl}_{\mathcal{M}, F_i, \hat{T}_i} := \text{ConcCl}(\text{cod}, N_{F_i}, \hat{T}_i)$. Similarly, taking the fibration for logical relations $\prod_{i \in \mathbb{I}} \text{cod}$ one obtains a category $\mathcal{K} := \text{Krip}(\prod_{i \in \mathbb{I}} \text{cod}, \langle N_{F_i} \rangle_{i \in \mathbb{I}})$ as well as subcategories $\mathcal{C} := \text{Conc}(\prod_{i \in \mathbb{I}} \text{cod}, \langle N_{F_i} \rangle_{i \in \mathbb{I}})$ and $\mathcal{C}Cl := \text{ConcCl}(\prod_{i \in \mathbb{I}} \text{cod}, \langle N_{F_i} \rangle_{i \in \mathbb{I}}, \hat{T})$, for $\hat{T}$ the monad defined in Lemma 7.1.

Because cartesian liftings are determined component-wise (Lemma 7.1(1)), so is $\hat{T}$ -closure. It follows that all the properties of Lemma 7.1 extend to $\mathcal{C}Cl$. Moreover, $\mathcal{C}Cl$ acquires a bi-CCC structure and a strong monad W with underlying functor $H\hat{T}k$ by Lemma 10.1. We call $(\mathcal{C}Cl, W, \hat{s})$ the *abstract OHR model* on $(\mathcal{M}, T, s)$. The situation is summarised below (cf. Figure 1b and diagram (11)):

$$\text{ConcCl}_{\mathcal{M}, F_i, \hat{T}_i} \xleftarrow{\perp} \text{Conc}_{\mathcal{M}, F_i} \xleftarrow{j} \text{Krip}_{\mathcal{M}, F_i} \quad W := H\hat{T}k \circlearrowleft \mathcal{C}Cl \xleftarrow{\perp} \mathcal{C} \xleftrightarrow[\perp]{\text{K}} \mathcal{K} \circlearrowright \hat{t}$$

Next we consider logical relations over $\mathcal{C}Cl$ (cf. diagram (7)). Extending Lemmas 7.2 and 7.3 to incorporate sums, one obtains the following extension of Prop. 7.1 with sum types.

Proposition 10.1. *Let $\mathcal{R} = \{\mathcal{R}_\sigma\}_{\sigma \in \text{Ty}^+}$ be a hungry, $(\lambda, +, 0)$ -compatible $\lambda_c^+(\mathbf{S})$ -logical relation over $\mathcal{W}^{\text{TT}[\mathcal{R}]}$ and suppose $i_0 \in \mathbb{I}$ is such that $\hat{T}_{i_0} = T^{\text{TT}[\langle \mathcal{R} \rangle]}$ and $\overline{\hat{s}[\sigma]}(i_0)(\beta) = \mathcal{R}_\beta^{\text{val}}$ for all $\beta \in \mathbf{B}$. Then $\overline{\hat{s}[\sigma]}(i_0) = \mathcal{R}_\sigma^{\text{val}}$ and $(\overline{\hat{W}\hat{s}[\sigma]})(i_0) = \mathcal{R}_\sigma$ for all $\sigma \in \text{Ty}^+$, so every $f : \hat{s}[\Gamma] \rightarrow \hat{W}\hat{s}[\sigma]$ in CCL satisfies $\mathcal{R}$.*

It remains to instantiate the abstract construction. The development goes through verbatim, except in the definition of $\mathbb{I}$ (equation (8)) we must now assume that $\text{Interp}(\mathbb{A}, F, \hat{T})$ consists of maps $r : \mathbf{B} \rightarrow \text{ConcCl}_{\mathcal{M}, F, \hat{T}}$ rather than $\mathbf{B} \rightarrow \text{Conc}_{\mathcal{M}, F}$. We therefore obtain an *OHR model* as in Def. 8.1. In this setting, Prop. 8.1 and Cor. 8.1 are as follows.

Theorem 10.1. *Let $\mathcal{R} = \{\mathcal{R}_\sigma\}_{\sigma \in \text{Ty}}$ be a hungry, $(\lambda, +, 0)$ -compatible $\lambda_c^+(\mathbf{S})$ -logical relation over $\hat{W} := \mathcal{W}^{\text{TT}[\mathcal{R}]}$ such that every global element $g : 1 \rightarrow s[\beta]$ in $\mathcal{M}$ is in $\mathcal{R}_\beta^{\text{val}}(\diamond)$. Then every $\text{OHR}(\mathcal{M})$ -morphism $f : \hat{s}[\Gamma] \rightarrow \hat{W}\hat{s}[\sigma]$ satisfies $\mathcal{R}$.*

Corollary 10.1. *If $F : (\mathcal{N}, S, t) \rightarrow (\text{OHR}(\mathcal{M}), W, \hat{s})$ is a $\lambda_c^+(\mathbf{S})$ -morphism with a surjective map $(U \circ F)_{1, t[\beta]} : \mathcal{N}(1, t[\beta]) \rightarrow \mathcal{M}(1, s[\beta])$, then any $f : \hat{s}[\Gamma] \rightarrow \hat{W}\hat{s}[\sigma]$ in $\text{OHR}(\mathcal{M})$ is in F 's image.*

Since Lemma 8.2 extends to sum types without difficulty, we also recover full abstraction.

Theorem 10.2. *Let $(\mathcal{M}, T, s)$ be a small, well-pointed $\lambda_c^+(\mathbf{S})$ -model such that for every $\beta \in \mathbf{B}$ and $g : 1 \rightarrow s[\beta]$ the composite $\eta_{s[\beta]} \circ g$ is definable. Then, if $\text{Krip}(\prod_i \text{cod}, \langle F \rangle_{i \in \mathbb{I}})$ admits a tractable hull functor and the counit of the adjunction $j \dashv H$ is component-wise monic, the induced model $(\text{OHR}(\mathcal{M}), W, \hat{s})$ is well-pointed and fully complete, hence fully abstract.*

11 EXAMPLES: FULLY ABSTRACT MODELS OVER Diff AND QBS

In this final section we give two examples of the OHR construction with sums. First we study a simple language for probabilistic programming. This is the fragment of the idealised Anglican language of Staton et al. [2016] without recursive types and constructors; for the semantics we use the category of *quasi-Borel spaces* [Heunen et al. 2017]. Then we turn to the language for automatic differentiation studied by Huot et al. [2020], extended with a global memory cell that may be read and written to. The semantics takes place in the category of *diffeological spaces* (e.g. Iglesias-Zemmour [2013]). In each case the hull functor is defined as in Remark 6.1, so tractable.

11.1 Probability over QBS

The category of *quasi-Borel spaces* was introduced by Heunen et al. [2017] as a setting for the denotational semantics of higher-order probabilistic programming languages. Naively, one would hope to interpret such languages in the category **Meas** of measurable spaces and measurable maps using the *Giry monad* [Giry 1982]. However, **Meas** is not cartesian closed; **QBS** rectifies this deficiency. **QBS** acts as a kind of conservative extension of the category of standard Borel spaces: there is a functor $R : \mathbf{Meas} \rightarrow \mathbf{QBS}$ and, if X, Y are standard Borel spaces, then $\mathbf{QBS}(RX, RY) = \mathbf{Meas}(X, Y)$ (see e.g. Ścibior et al. [2018]). Moreover, the Giry monad restricts to a monad P on **QBS**, and if X, Y are standard Borel spaces then maps $RX \rightarrow P(RY)$ correspond to the *s-finite kernels* used by Staton [2017] to give a complete semantics for a first-order probabilistic language.

The signature S_{prob} for our idealised probabilistic programming language is given in the box below. The type `real` represents the real numbers. To get a primitive representing each measurable function we use the (co)cartesian structure of **Meas**: where $m : \{\text{real}\} \rightarrow \mathbf{Meas}$ interprets `real` as the standard Borel space $(\mathbb{R}, \Sigma_{\mathbb{R}})$, one obtains a measurable space $m[\gamma]$ for all ground types γ . Constant maps and all the usual distributions are measurable; hence for any $r \in \mathbb{R}$ one has $\text{const}_r : 1 \rightarrow \text{real}$.

base types: `real`;
primitives: $f : \gamma \rightarrow \nu$ for every $f : m[\gamma] \rightarrow m[\nu]$;
operations (for $\gamma \in \mathbf{G}$): `score` : $\text{real} \rightsquigarrow 1$, `sample $\gamma$` : $\gamma \rightsquigarrow \gamma$, `normalise $\gamma$` : $\gamma \rightsquigarrow \text{real} * \gamma + 1 + 1$;

Semantically, we set $s(\text{real})$ to be the quasi-Borel space corresponding to $(\mathbb{R}, \Sigma_{\mathbb{R}})$. Because $\mathbb{R}$ is standard Borel, $m[\gamma]$ is standard Borel for every ground type γ . For each operation $\text{op} : \gamma \rightsquigarrow \nu$ we therefore take $s[\text{op}]$ to be the Kleisli arrow corresponding to the s -finite kernel interpretation of that operation given by Staton [2017]. Finally, for a primitive $\bar{f} : \gamma \rightarrow \nu$ we use the fact that $\text{QBS}(s[\gamma], s[\nu]) = \text{Meas}(m[\gamma], m[\nu])$ to define $s[\bar{f}] := \lambda(* \in 1) . \bar{f} : 1 \rightarrow (s[\gamma] \Rightarrow s[\nu])$. Thus, we have a $\lambda_c^+(\text{S}_{\text{prob}})$ -model (QBS, P, s) (we silently identify QBS with a suitable small subcategory).

The induced OHR model is fully abstract, since for every $g : 1 \rightarrow s[\text{real}]$ the composite $\eta_{s[\text{real}]}^P \circ g$ is definable. Indeed, global elements in QBS are in bijective correspondence with global elements in Set , so we can identify g with a constant map const_r for some $r \in \mathbb{R}$. Then $\eta_{s[\text{real}]}^P \circ g = s[\diamond \vdash \text{const}_r () : \text{real}]$, as required. So the OHR model over (QBS, P, s) exists and is fully abstract.

The forgetful functor $\text{QBS} \rightarrow \text{Set}$ strictly preserves products and coproducts and evaluation in QBS is as in Set , so the relationship between $\text{OHR}(\text{QBS})$ and QBS is very similar to that between $\text{OHR}(\text{Fin})$ and Fin outlined in Sec. 9. The monad W is a restriction of P , products in $\text{OHR}(\text{QBS})$ are as in QBS —hence as in Set —and the evaluation map is a restriction of that in Set (cf. (9)).

11.2 Global State over Diff

Just as quasi-Borel spaces were introduced to deal with the fact that Meas is not cartesian closed, so the category Diff of diffeological spaces and smooth maps was used by Huot et al. [2020] to deal with the fact that the usual setting for differential geometry, namely the category of cartesian spaces and smooth functions, is not cartesian closed. A *diffeological space* is a pair $(X, \mathcal{P}_X)$ consisting of a set X equipped with a set of *plots* $\mathcal{P}_X^U \subseteq \text{Set}(U, X)$ for every $n \in \mathbb{N}$ and open subset $U \subseteq \mathbb{R}^n$, subject to certain axioms. One then widens the definition of smooth function from differential geometry to this clean axiomatic setting: one calls a function *smooth* if it send plots to plots.

Huot et al. [2020] show that Diff is a natural setting for studying the denotational semantics of syntactic (forward mode) automatic differentiation for neural network programming. We consider an extension of their simple language for AD with a single memory cell. One may *lookup* the value of the cell or *update* it to a new shared value; we assume values range over a fixed set V . The relevant signature S_{GS} is defined in the box below, where the primitives $+$ and $\times$ represent the usual operations on $\mathbb{R}$ and ς represents the sigmoid function $\varsigma(x) = (1 + e^{-x})^{-1}$.

base types: real, Val ; **operations:** $\text{lookup} : 1 \rightsquigarrow \text{Val}$; $\text{update} : \text{Val} \rightsquigarrow 1$;
primitives: $\underline{r} : \text{real}$ (for $r \in \mathbb{R}$); $+, \times : \text{real} * \text{real} \rightarrow \text{real}$; $\varsigma : \text{real} \rightarrow \text{real}$; $\underline{v} : \text{Val}$ (for $v \in V$).

For the semantic interpretation, set $s(\text{real}) := (\mathbb{R}, \mathcal{P}_{\mathbb{R}})$, where $\mathcal{P}_{\mathbb{R}}^U$ is the set of smooth maps $U \rightarrow \mathbb{R}$. We interpret Val as a *coarse* diffeological space: $s(\text{Val}) := (V, \mathcal{D}_V)$ where $\mathcal{D}_V^U := \text{Set}(U, V)$. The memory cell is modelled by the *global state monad* $G_V(X) := V \Rightarrow (V \times X)$ on Diff . Primitives are interpreted by the corresponding (smooth) functions in Set ; the usual set maps interpreting lookup and update (e.g. [Kammar 2014]) are smooth because $s(\text{Val})$ is coarse. This defines a $\lambda_c^+(\text{S}_{\text{GS}})$ -model (Diff, G_V, s) . The global elements in Diff are exactly the global elements in Set so, identifying Diff with a suitable small subcategory, the OHR model over (Diff, G_S, s) exists and is fully abstract.

The forgetful functor $\text{Diff} \rightarrow \text{Set}$ strictly preserves products and coproducts, and evaluation in Diff is evaluation in Set . Thus, just as for the QBS example, the relationship between the OHR model and (Diff, G_S, s) is very similar to that between $\text{OHR}(\text{Fin})$ and Fin outlined in Sec. 9.

ACKNOWLEDGMENTS

We thank the authors of [Matache et al. 2021] for helpful discussions relating their work to ours, C. Matache & S. Staton for the models in Exs. 1.1 and 1.2, and D. McDermott, H. Paquet, N. Arkor and P. B. Levy for insightful comments. Finally we thank the reviewers for their thoughtful suggestions.

REFERENCES

- S. Abramsky, K. Honda, and G. McCusker. 1998. A fully abstract game semantics for general references. In *Proceedings of the Thirteenth Annual IEEE Symposium on Logic in Computer Science (Cat. No.98CB36226)*. IEEE Comput. Soc. <https://doi.org/10.1109/lics.1998.705669>
- S. Abramsky and R. Jagadeesan. 1994. Games and full completeness for multiplicative linear logic. *Journal of Symbolic Logic* 59, 2 (June 1994), 543–574. <https://doi.org/10.2307/2275407>
- S. Abramsky, R. Jagadeesan, and P. Malacaria. 2000. Full Abstraction for PCF. *Information and Computation* 163, 2 (Dec. 2000), 409–470. <https://doi.org/10.1006/inco.2000.2930>
- J. Adamek, H. Herrlich, and G.E. Strecker. 2009. *Abstract and Concrete Categories: The Joy of Cats*. Dover Publications. <https://books.google.co.jp/books?id=rqT4PgAACAAJ>
- M. Alimohamed. 1995. A Characterization of Lambda Definability in Categorical Models of Implicit Polymorphism. *Theor. Comput. Sci.* 146, 1-2 (July 1995), 5–23. [https://doi.org/10.1016/0304-3975\(94\)00283-O](https://doi.org/10.1016/0304-3975(94)00283-O)
- F. Borceux. 1994. *Handbook of Categorical Algebra, volume 1*. Cambridge University Press. <https://doi.org/10.1017/cbo9780511525858>
- R. Cartwright, P.L. Curien, and M. Felleisen. 1994. Fully Abstract Semantics for Observably Sequential Languages. 111, 2 (jun 1994), 297–401. <https://doi.org/10.1006/inco.1994.1047>
- P. Clairambault and M. de Visme. 2020. Full abstraction for the quantum lambda-calculus. *Proceedings of the ACM on Programming Languages* 4, POPL (jan 2020), 1–28. <https://doi.org/10.1145/3371131>
- P.-L. Curien. 2007. Definability and Full Abstraction. *Electronic Notes in Theoretical Computer Science* 172 (April 2007), 301–310. <https://doi.org/10.1016/j.entcs.2007.02.011>
- U. de’ Liguoro. 1996. *PCF Definability via Kripke Logical Relations (after O’Hearn and Riecke)*. Technical Report. Laboratoire d’Informatique de l’Ecole Normale Supérieure.
- T. Ehrhard, C. Tasson, and M. Pagani. 2014. Probabilistic coherence spaces are fully abstract for probabilistic PCF. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM. <https://doi.org/10.1145/2535838.2535865>
- M. Fiore, A. Jung, E. Moggi, P. O’Hearn, Riecke J., G. Rosolini, and I. Stark. 1996. Domains and denotational semantics: History, accomplishments and open problems. *Bulletin of EATCS* 59 (1996), 227–256. Edited by A. Jung.
- M. Fiore, G. Plotkin, and D. Turi. 1999. Abstract Syntax and Variable Binding. In *Proceedings of the 14th Annual IEEE Symposium on Logic in Computer Science (LICS ’99)*. IEEE Computer Society, Washington, DC, USA, 193–. <http://dl.acm.org/citation.cfm?id=788021.788948>
- M. Fiore and A. Simpson. 1999. Lambda Definability with Sums via Grothendieck Logical Relations. In *Typed Lambda Calculi and Applications*, J.-Y. Girard (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 147–161.
- P. Freyd. 1972. Aspects of topoi. *Bulletin of the Australian Mathematical Society* 7, 1 (aug 1972), 1–76. <https://doi.org/10.1017/s0004972700044828>
- J. Y. Girard. 1989. *Proofs and types*. Cambridge University Press, Cambridge.
- M. Giry. 1982. A categorical approach to probability theory. In *Lecture Notes in Mathematics*. Springer Berlin Heidelberg, 68–85. <https://doi.org/10.1007/bfb0092872>
- J. Goubault-Larrecq, S. Lasota, and D. Nowak. 2008. Logical relations for monadic types. *Mathematical Structures in Computer Science* 18, 06 (Oct. 2008), 1169. <https://doi.org/10.1017/s0960129508007172>
- J. Goubault-Larrecq, S. Lasota, D. Nowak, and Y. Zhang. 2004. Complete Lax Logical Relations for Cryptographic Lambda-Calculi. In *Computer Science Logic*. Springer Berlin Heidelberg, 400–414. https://doi.org/10.1007/978-3-540-30124-0_31
- C. A. Hermida. 1993. *Fibrations, Logical Predicates and Indeterminates*. Ph.D. Dissertation. University of Edinburgh.
- C. Heunen, O. Kammar, S. Staton, and H. Yang. 2017. A Convenient Category for Higher-Order Probability Theory. In *Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (Reykjavík, Iceland) (LICS ’17)*. IEEE Press, Article 77, 12 pages.
- M. Huot, S. Staton, and M. Vákár. 2020. Correctness of Automatic Differentiation via Diffeologies and Categorical Gluing. In *Lecture Notes in Computer Science*. Springer International Publishing, 319–338. https://doi.org/10.1007/978-3-030-45231-5_17
- J.M.E. Hyland and C.-H.L. Ong. 2000. On Full Abstraction for PCF: I, II, and III. *Information and Computation* 163, 2 (Dec. 2000), 285–408. <https://doi.org/10.1006/inco.2000.2917>
- P. Iglesias-Zemmour. 2013. *Diffeology*. American Mathematical Society, Providence, Rhode Island.
- B. Jacobs. 1993. Comprehension categories and the semantics of type dependency. *Theoretical Computer Science* 107, 2 (Jan. 1993), 169–207. [https://doi.org/10.1016/0304-3975\(93\)90169-t](https://doi.org/10.1016/0304-3975(93)90169-t)
- B. Jacobs. 1999. *Categorical Logic and Type Theory*. Number 141 in Studies in Logic and the Foundations of Mathematics. North Holland, Amsterdam.
- P. T. Johnstone. 2002. *Sketches of an elephant : a topos theory compendium*. Oxford University Press, Oxford New York.

- A. Jung and J. Tiuryn. 1993. A new characterization of lambda definability. In *Typed Lambda Calculi and Applications*, Marc Bezem and Jan Friso Groote (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 245–257.
- O. Kammar. 2014. *An Algebraic Theory of Type-and-Effect Systems*. Ph.D. Dissertation. University of Edinburgh.
- O. Kammar and D. McDermott. 2018. Factorisation Systems for Logical Relations and Monadic Lifting in Type-and-effect System Semantics. *Electronic Notes in Theoretical Computer Science* 341 (2018), 239 – 260. <https://doi.org/10.1016/j.entcs.2018.11.012> Proceedings of the Thirty-Fourth Conference on the Mathematical Foundations of Programming Semantics (MFPS XXXIV).
- O. Kammar and G. D. Plotkin. 2012. Algebraic foundations for effect-dependent optimisations. In *Proceedings of the 39th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages - POPL '12*. ACM Press. <https://doi.org/10.1145/2103656.2103698>
- S. Katsumata. 2005. A Semantic Formulation of $\top\top$ -Lifting and Logical Predicates for Computational Metalanguage. In *Computer Science Logic*. Springer Berlin Heidelberg, 87–102. https://doi.org/10.1007/11538363_8
- S. Katsumata. 2008. A Characterisation of Lambda Definability with Sums Via $\top\top$ -Closure Operators. In *Computer Science Logic*, Michael Kaminski and Simone Martini (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 278–292.
- S. Katsumata. 2013. Relating computational effects by $\top\top$ -lifting. *Information and Computation* 222 (2013), 228 – 246. <https://doi.org/10.1016/j.ic.2012.10.014> 38th International Colloquium on Automata, Languages and Programming (ICALP 2011).
- A. Kock. 1972. Strong functors and monoidal monads. *Archiv der Mathematik* 23, 1 (dec 1972), 113–120. <https://doi.org/10.1007/bf01304852>
- S. Lasota, D. Nowak, and Y. Zhang. 2007. On Completeness of Logical Relations for Monadic Types. In *Advances in Computer Science - ASIAN 2006. Secure Software and Related Issues*. Springer Berlin Heidelberg, 223–230. https://doi.org/10.1007/978-3-540-77505-8_17
- F. W. Lawvere. 2006. Diagonal Arguments and Cartesian Closed Categories. *Reprints in Theory and Applications of Categories* 15 (2006), 1–13. <http://www.tac.mta.ca/tac/reprints/articles/15/tr15.pdf>
- F. Loregian and E. Riehl. 2020. Categorical notions of fibration. *Expositiones Mathematicae* 38, 4 (Dec. 2020), 496–514. <https://doi.org/10.1016/j.exmath.2019.02.004>
- Q. M. Ma and John C. Reynolds. 1992. Types, abstraction, and parametric polymorphism, part 2. In *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 1–40. https://doi.org/10.1007/3-540-55511-0_1
- M. Marz. 2000. A fully abstract model for sequential computation. *Electronic Notes in Theoretical Computer Science* 35 (2000), 133–152. [https://doi.org/10.1016/s1571-0661\(05\)80735-2](https://doi.org/10.1016/s1571-0661(05)80735-2)
- C. Matache, S. Moss, , and S. Staton. 2021. Recursion and Sequentiality in Categories of Sheaves. *To appear in Proceedings of 6th International Conference on Formal Structures for Computation and Deduction (FSCD 2021)* (2021). <https://doi.org/10.4230/LIPIcs.FSCD.2021.25>
- R. Milner. 1977. Fully abstract models of typed λ -calculi. *Theoretical Computer Science* 4, 1 (Feb. 1977), 1–22. [https://doi.org/10.1016/0304-3975\(77\)90053-6](https://doi.org/10.1016/0304-3975(77)90053-6)
- J. C. Mitchell and A. Scedrov. 1993. Notes on sconing and relators. In *Computer Science Logic*. Springer Berlin Heidelberg, 352–378. https://doi.org/10.1007/3-540-56992-8_21
- E. Moggi. 1989. Computational Lambda-Calculus and Monads. In *Proceedings of the Fourth Annual Symposium on Logic in Computer Science* (Pacific Grove, California, USA). IEEE Press, 14–23.
- E. Moggi. 1991. Notions of computation and monads. *Inf. Comput.* 93, 1 (1991), 55–92.
- A. S. Murawski and N. Tzevelekos. 2012. Algorithmic Games for Full Ground References. In *Automata, Languages, and Programming*. Springer Berlin Heidelberg, 312–324. https://doi.org/10.1007/978-3-642-31585-5_30
- P. W. O’Hearn and J. G. Riecke. 1995. Kripke Logical Relations and PCF. *Information and Computation* 120, 1 (1995), 107 – 116. <https://doi.org/10.1006/inco.1995.1103>
- B. Pareigis. 1977. Non-additive ring and module theory II: C-categories, C-functors and C-morphisms. *Publicationes mathematicae* 24, 3 (January 1977).
- G. D. Plotkin. 1973. *Lambda-definability and logical relations*. Technical Report. University of Edinburgh.
- G. D. Plotkin. 1977. LCF considered as a programming language. *Theoretical Computer Science* 5, 3 (Dec. 1977), 223–255. [https://doi.org/10.1016/0304-3975\(77\)90044-5](https://doi.org/10.1016/0304-3975(77)90044-5)
- G. D. Plotkin. 1980. Lambda-Definability in the Full Type Hierarchy. In *To H. B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*, Jonathan P. Seldin and J. Roger Hindley (Eds.). Academic Press.
- G. D. Plotkin and J. Power. 2003. Algebraic Operations and Generic Effects. *Applied Categorical Structures* 11, 1 (2003), 69–94. <https://doi.org/10.1023/a:1023064908962>
- J. Power and E. Robinson. 2000. Logical Relations and Data Abstraction. In *Computer Science Logic*. Springer Berlin Heidelberg, 497–511. https://doi.org/10.1007/3-540-44622-2_34
- J. G. Riecke and A. Sandholm. 2002. A Relational Account of Call-by-Value Sequentiality. *Information and Computation* 179, 2 (Dec. 2002), 296–331. <https://doi.org/10.1006/inco.2002.2957>

- G. Scherer. 2017. Deciding Equivalence with Sums and the Empty Type. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages* (Paris, France) (POPL 2017). Association for Computing Machinery, New York, NY, USA, 374–386. <https://doi.org/10.1145/3009837.3009901>
- A. Ścibior, O. Kammar, V. Vákár, S. Staton, H. Yang, Y. Cai, K. Ostermann, S. K. Moss, C. Heunen, and Z. Ghahramani. 2018. Denotational validation of higher-order Bayesian inference. *Proceedings of the ACM on Programming Languages* 2, POPL (Jan. 2018), 1–29. <https://doi.org/10.1145/3158148>
- J. M. Souriau. 1980. Groupes différentiels. In *Differential Geometrical Methods in Mathematical Physics*, P. L. García, A. Pérez-Rendón, and J. M. Souriau (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 91–128.
- S. Staton. 2017. Commutative Semantics for Probabilistic Programming. In *Programming Languages and Systems*. Springer Berlin Heidelberg, 855–879. https://doi.org/10.1007/978-3-662-54434-1_32
- S. Staton, H. Yang, F. Wood, C. Heunen, and O. Kammar. 2016. Semantics for probabilistic programming. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science*. ACM. <https://doi.org/10.1145/2933575.2935313>
- R. Street. 1972. The formal theory of monads. *Journal of Pure and Applied Algebra* 2, 2 (1972), 149–168. [https://doi.org/10.1016/0022-4049\(72\)90019-9](https://doi.org/10.1016/0022-4049(72)90019-9)