

Effectful Applicative Bisimilarity: Monads, Relators, and Howe's Method

Ugo Dal Lago, Francesco Gavazzo, Paul Blain Levy

► **To cite this version:**

Ugo Dal Lago, Francesco Gavazzo, Paul Blain Levy. Effectful Applicative Bisimilarity: Monads, Relators, and Howe's Method. LICS 2017 - ACM/IEEE Symposium on Logic in Computer Science, Jun 2017, Reykjavik, Iceland. IEEE, Logic in Computer Science (LICS), 2017 32nd Annual ACM/IEEE Symposium on, pp.1-12, <10.1109/LICS.2017.8005117>. <hal-01636365>

HAL Id: hal-01636365

<https://hal.inria.fr/hal-01636365>

Submitted on 5 Dec 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Effectful Applicative Bisimilarity: Monads, Relators, and Howe’s Method

Ugo Dal Lago and Francesco Gavazzo
Università di Bologna & INRIA Sophia Antipolis
Bologna, Italy

Paul Blain Levy
University of Birmingham
Birmingham, United Kingdom

Abstract—We study Abramsky’s applicative bisimilarity abstractly, in the context of call-by-value λ -calculi with algebraic effects. We first of all endow a computational λ -calculus with a monadic operational semantics. We then show how the theory of relators provides precisely what is needed to generalise applicative bisimilarity to such a calculus, and to single out those monads and relators for which applicative bisimilarity is a congruence, thus a sound methodology for program equivalence. This is done by studying Howe’s method in the abstract.

I. INTRODUCTION

Program equivalence is one of the central notions in the theory of programming languages, and giving satisfactory definitions and methodologies for it is a challenging problem, for example when dealing with higher-order languages. The problem has been approached, since the birth of the discipline, in many different ways. One can define program equivalence through denotational semantics, thus relying on a model and stipulating two programs to be equivalent if and only if they are interpreted by the same denotation. If the calculus at hand is equipped with a notion of *observation*, typically given through some forms of operational semantics, one could proceed following the route traced by Morris [31], and define programs to be *contextually* equivalent when they *behave* the same *in every* context.

Both these approaches have their drawbacks, the first one relying on the existence of a (not too coarse) denotational model, the latter quantifying over all contexts, and thus making concrete proofs of equivalence hard. Handier methodologies for proving programs equivalent have been introduced along the years based on logical relations and applicative bisimilarity. Logical relations were originally devised for typed, normalising languages, but later generalised to more expressive formalisms, e.g., through step-indexing [2] and biorthogonality [5]. Starting from Abramsky’s pioneering work on applicative bisimilarity [1], coinduction has also been proved to be a useful methodology for program equivalence, and has been applied to a variety of calculi and language features.

The scenario just described also holds when the underlying calculus is not pure, but effectful. There have been many attempts to study effectful λ -calculi [34], [30] by way of denotational semantics [21], [14], [12], logical relations [6], and applicative bisimilarity [25], [10], [8]. But while the denotational and logical relation semantics of effectful calculi have been studied in the abstract [18], [20], the same cannot be said about applicative bisimilarity and related coinductive

techniques. There is a growing body of literature on applicative bisimilarity for calculi with, e.g., nondeterministic [25], and probabilistic effects [10], but each notion of an effect has been studied independently, often getting different results. Distinct proofs of congruence for applicative bisimilarity, even if done through a common methodology, namely the so-called Howe’s method [19], do not at all have the same difficulty in each of the cases cited above. As an example, the proof of the so-called Key Lemma relies on duality results from linear programming [38] when done for probabilistic effects, contrarily to the apparently similar case of nondeterministic effects, whose logical complexity is comparable to that for the plain, deterministic λ -calculus [32], [25]. Finally, as the third author observed in his work with Koutavas and Sumii [23], applicative bisimilarity is fragile to the presence of certain effects, like local states or dynamically created exceptions: in these cases, a sort of information hiding is possible which makes applicative bisimilarity simply too weak, and thus unsound for contextual equivalence.

The observations above naturally lead to some questions. Is there any way to factor out the common part of the congruence proof for applicative bisimilarity in the cases above? Where do the limits on the correctness of applicative bisimilarity lie, in presence of effects? The authors strongly believe that the field of coinductive techniques for higher-order program equivalence should be better understood *in the abstract*, this way providing some answers to the questions above, given that generic accounts for effectful λ -calculi abound in the literature [30], [34].

This paper represents a first step towards answering the questions above. We first of all introduce a computational λ -calculus in which general algebraic effects can be represented, and give a monadic operational semantics for it, showing how the latter coincides with the expected one in many distinct concrete examples. We then show how applicative bisimilarity can be defined for any instance of such a monadic λ -calculus, based on the notion of a relator, which allows to account for the possible ways a relation on a set X can be turned into one for TX , where T is a monad. We then single out a set of axioms for monads and relators which allow us to follow Howe’s proof of congruence for applicative bisimilarity *in the abstract*. Noticeably, these axioms are satisfied in all the example algebraic effects we consider. The proof of it allows us to understand the deep reasons *why*, say, different instances of Howe’s method in the literature seem to have

$$\begin{aligned}
W &\rightarrow V \oplus \text{COMP}(V, W) \\
Z &\rightarrow T \underline{1} \\
T \underline{n} &\rightarrow (R \underline{n}) \oplus (T \underline{n+1}) \\
R \underline{0} &\rightarrow \lambda x.x \\
R \underline{n+1} &\rightarrow \text{COMP}(R \underline{n}, V)
\end{aligned}$$

Fig. 1. Two Probabilistic Programs.

different complexities.

Due to space constraints, many proofs need to be elided. An extended version of this paper with more details is available [9].

II. ON COINDUCTION AND EFFECTFUL λ -CALCULI

In this section, we illustrate how coinduction can be useful when proving the equivalence of programs written in higher-order effectful calculi.

Let us start with a simple example of two supposedly equivalent probabilistic functional programs, W and Z , given in Figure 1. (The expression $\text{COMP}(M, N)$ stands for the term $\lambda y.M(Ny)$, and \oplus is a binary operation for fair probabilistic choice.) Both W and Z behave like the n -th composition of a function V with itself with probability $\frac{1}{2^n}$, for every $n > 0$. But how could we even *define* the equivalence of such effectful programs? A natural answer consists in following Morris [31], and stipulate that two programs are contextually equivalent if they behave the same when put in any context, where the observable behaviour of a term can be taken, e.g., as its probability of convergence. Proving two terms to be contextually equivalent can be quite hard, given the universal quantification over all contexts on which contextual equivalence is based.

Applicative bisimilarity is an alternative definition of program equivalence, in which λ -terms are seen as computational objects interacting with their environment by exposing their behaviour, and by taking arguments as input. Applicative bisimilarity has been generalised to effectful λ -calculi of various kinds, and in particular to untyped probabilistic λ -calculi [10]. It is known to be not only a congruence (thus *sound* for contextual equivalence) but also *fully abstract*, at least for call-by-value evaluation [8]. Indeed, applicative bisimilarity can be applied to the example terms in Figure 1, which can this way be proved contextual equivalent (see [9] for the details).

The proof of soundness of applicative bisimilarity in presence of probabilistic effects is significantly more complicated than the original one, although both can be done by following the so-called Howe’s method [19]. More specifically, the proof that the Howe extension of similarity is a simulation relies on duality from linear programming (through the Max Flow Min Cut Theorem) when done in presence of probabilistic effects,

$$\begin{aligned}
W^{\text{raise}} &\rightarrow (V \oplus \text{raise}_e) \oplus \text{COMP}(V, W^{\text{raise}}) \\
Z^{\text{raise}} &\rightarrow T \underline{1} \\
T \underline{n} &\rightarrow ((R \underline{n}) \oplus \text{raise}_e) \oplus (T \underline{n+1}) \\
R \underline{0} &\rightarrow \lambda x.x \\
R \underline{n+1} &\rightarrow \text{COMP}(R \underline{n}, V)
\end{aligned}$$

Fig. 2. Two Probabilistic Programs Throwing Exceptions.

something that is not required in the plain, deterministic setting, nor in presence of *nondeterministic* choice.

Modern functional programming languages, however, can be “effectful” in quite complex ways. As an example, programs might be allowed not only to evolve probabilistically, but also to have an internal state, to throw exceptions, or to perform some input-output operations. Consider, as another simple example, the programs in Figure 2, a variation on the programs from Figure 1 where we allow an exception e to be thrown by way of the raise_e command. Intuitively, W^{raise} and Z^{raise} behave like W and Z , respectively, but they both throw an exception with a certain probability.

While applicative similarity in presence of dynamically created exceptions is well-known to be unsound [23], the mere presence of the raise_e command does not seem to cause any significant problem. The literature, however, does not offer any result about whether *combining* two or more notions of computational effect for which bisimilarity is known to work well, should be problematic or not. An *abstract* theory accounting for how congruence proofs can be carried out in effectful calculi is simply lacking.

Even if staying within the scope of Howe’s method, it seems that each effect between those analysed in the literature is handled by way of some *ad-hoc* notion of bisimulation. As an example, nondeterministic extensions of the λ -calculus can be dealt with by looking at terms as a labelled transition system, while probabilistic extensions of the λ -calculus require a different definition akin to Larsen and Skou’s probabilistic bisimulation [10]. What kind of transition system do we need when, e.g., dealing with the example programs from Figure 2? In other words, an abstract theory of effectful applicative bisimilarity would be beneficial from a purely definitional viewpoint, too.

What could come to the rescue here is the analysis of effects and bisimulation which has been carried out in the field of coalgebra [36]. In particular, we here exploit the theory of relators, also known as lax extensions [4], [39].

III. DOMAINS AND MONADS: SOME PRELIMINARIES

In this section, we recall some basic definitions and results on complete partial orders, categories, and monads. Those will be central in the rest of this paper but, due to space constraints, there is no hope to be comprehensive. We refer to this paper’s long version [9] or to the many introductory textbooks on

partial order theory [13] or category theory [29] for more details.

A. Domains and Continuous Σ -algebras

We assume basic familiarity with domain theory. In particular, we assume the notions of ω -complete partial order (ω CPO for short), and of pointed ω -complete partial order (ω CPPO for short), as well as the notions of monotone, continuous, and strict functions, which are standard. Finally, we recall that the category of ω CPPOs and continuous function is cartesian closed, meaning in particular that ω CPPOs are closed under (finite) products and exponentials, defined as continuous functions spaces. Following [34], [35], we consider operations (like \oplus or raise_e in the examples from Section II) from a given signature as sources of effects. Semantically, dealing with operation symbols requires the introduction of appropriate algebraic structures interpreting such operation symbols as suitable functions. Combining the algebraic and the order theoretic structures just described, leads to consider algebras carrying a domain structure (ω CPPO, in this paper), such that *all* function symbols are interpreted as continuous functions. The formal notion capturing all these desiderata is the one of a continuous Σ -algebra [15].

Recall that a signature $\Sigma = (\mathcal{F}, \alpha)$ consists of a set \mathcal{F} of operation symbols and a map $\alpha : \mathcal{F} \rightarrow \mathbb{N}$, assigning to each operation symbol a (finite) arity.

Definition 1. *Given a signature Σ , a continuous Σ -algebra is an ω CPPO (D, \sqsubseteq, \perp) such that for any function symbol σ in Σ there is an associated continuous function $\sigma^D : D^{\alpha(\sigma)} \rightarrow D$.*

Please observe that for a function symbol $\sigma \in \Sigma$, we do not require σ^D to be strict.

Before looking at monads, we now give various examples of concrete algebras which can be given the structure of a continuous Σ -algebra for certain signatures. This testifies the applicability of our theory to a relatively wide range of effects.

Example 1. *Let X be a set: the following are examples of ω CPPO.*

- *The flat lifting X_\perp of X , defined as $X + \{\perp\}$, ordered as follows: $x \sqsubseteq y$ iff $x = \perp$ or $x = y$.*
- *The set $(X + E)_\perp$ (think to E as a set of exceptions), ordered as in the previous example. We can consider the signature $\Sigma = \{\text{raise}_e \mid e \in E\}$, where each operation symbol raise_e is interpreted as the constant $\text{in}_l(\text{in}_r(e))$.*
- *The powerset $\mathcal{P}X$, ordered by inclusion. The least upper bound of a chain of sets is their union, whereas the bottom is the empty set. We can consider the signature $\Sigma = \{\oplus\}$ containing a binary operation symbol for nondeterministic choice. The latter can be interpreted as (binary) union, which is indeed continuous.*
- *The set of subdistributions $\mathcal{D}X = \{\mu : X \rightarrow [0, 1] \mid \text{supp}(\mu) \text{ countable, } \sum_{x \in X} \mu(x) \leq 1\}$ over X , ordered pointwise: $\mu \sqsubseteq \nu$ iff $\forall x \in X. \mu(x) \leq \nu(x)$. Note that requiring the support of μ to be countable is equivalent*

to requiring the existence of $\sum_{x \in X} \mu(x)$. The ω CPPO structure is pointwise induced by the one of $[0, 1]$ with the natural ordering. The least element is the always zero distribution $x \mapsto 0$ (note that the latter is a subdistribution, and not a distribution). We can consider the signature $\Sigma = \{\oplus_p \mid p \in [0, 1]\}$ with a family of probabilistic choice operations indexed by real numbers in $[0, 1]$. We can interpret \oplus_p as the binary operation $(x, y) \mapsto p \cdot x + (1-p) \cdot y$, which is indeed continuous.

- *The set $(S \times X)_\perp^S$, or equivalently $S \multimap (X \times S)$ (the set of partial states over X) with extension order: $f \sqsubseteq g$ iff $\forall x \in X. f(x) \neq \perp \Rightarrow f(x) = g(x)$, for a fixed set S (of states). The bottom element is the totally undefined function $x \mapsto \perp$, whereas the least upper bound of a chain $(f_n)_{n < \omega}$ is computed pointwise. Depending on the choice of S , we can define several continuous operations on $(S \times X)_\perp^S$. For instance, taking $S = \{\text{true}, \text{false}\}$, the set of booleans, we can consider the signature $\Sigma = \{\text{read}, \text{write}_b \mid b \in S\}$ to be interpreted as the continuous operations *read* and *write_b* defined by*

$$\text{write}_b(f) = x \mapsto f(b);$$

$$\text{read}(f, g) = x \mapsto \text{if } x = \text{true} \text{ then } f(x) \text{ else } g(x).$$

B. Monads

The notion of monad is given via the equivalent notion of Kleisli Triple (see [29]). Let \mathbb{C} be a category.

Definition 2. *A Kleisli Triple $\langle T, \eta, (\cdot)^\dagger \rangle$ consists of an endomap T over objects of \mathbb{C} , a family of arrows η_X , for any object X , and an operation (called Kleisli extension or Kleisli star) $(\cdot)^\dagger : \text{Hom}_{\mathbb{C}}(X, TY) \rightarrow \text{Hom}_{\mathbb{C}}(TX, TY)$, (for all objects X, Y) such that $f^\dagger \circ \eta = f$, $\eta^\dagger = \text{id}$ and $(g^\dagger \circ f)^\dagger = g^\dagger \circ f^\dagger$ hold, for f and g of the appropriate types.*

We will often denote a Kleisli Triple $\langle T, \eta, (\cdot)^\dagger \rangle$ simply as T . From now on we fix the base category \mathbb{C} to be the category \mathbb{SET} of sets and functions.

Remark 1. *Since we work in \mathbb{SET} , we will extensively use the so called bind operator $\gg=$ in place of Kleisli extensions. Such operator takes as arguments an element u of TX , together with a function $f : X \rightarrow TY$ and returns an element $u \gg= f$ in TY . Concretely, we can define $u \gg= f$ as $f^\dagger(u)$. Vice versa, we can define the Kleisli extension f^\dagger of f as $x \mapsto (x \gg= f)$.*

Example 2. *All the constructions introduced in Example 1 are monadic. For instance¹:*

- *The functor $TX = X_\perp$ is (part of) a monad, with left injection as unit and bind operator defined by*

$$u \gg= f = \begin{cases} f(x) & \text{if } u = \text{in}_l(x), \text{ for some } x \in X; \\ \text{in}_r(\perp) & \text{otherwise.} \end{cases}$$

¹We will be terminologically sloppy, not distinguishing between monads and Kleisli Triples not even on a terminological level. For instance, in the following example we will often speak of functors and monads, although what we really mean is endomaps on objects (of \mathbb{SET}) and Kleisli triples.

- The powerset functor \mathcal{P} is a monad with unit $x \mapsto \{x\}$ and bind operator defined by $u \gg f = \bigcup_{x \in u} f(x)$.
- The subdistribution functor \mathcal{D} is a monad with unit given via the Dirac distribution δ and bind operator defined by

$$\mu \gg f = y \mapsto \sum_{x \in X} \mu(x) \cdot f(x)(y).$$

- The partiality and exception functor $TX = (X + E)_\perp$ for a given set E of exceptions is a monad with the function $x \mapsto \text{in}_l(\text{in}_l(x))$ as unit. The bind operator is defined by

$$u \gg f = \begin{cases} u & \text{if } u = \text{in}_r(\perp) \text{ or } u = \text{in}_l(\text{in}_r(e)); \\ f(x) & \text{if } u = \text{in}_l(\text{in}_l(x)). \end{cases}$$

- The partiality and global state functor $TX = S \rightarrow (X \times S)_\perp$ for a given set S of states, is a monad with unit $x \mapsto (s \mapsto (x, s))$ and the bind operator defined by

$$(\sigma \gg f)(s) = \begin{cases} \perp & \text{if } \sigma(s) = \perp; \\ f(y)(t) & \text{if } \sigma(s) = (y, t). \end{cases}$$

For a given signature Σ , we are interested in monads on \mathbb{SET} that carry a continuous Σ -algebra structure.

Definition 3. An ωCPPO order \sqsubseteq on a monad T is a map that assigns to each set X a relation $\sqsubseteq_X \subseteq TX \times TX$ and an element $\perp_X \in TX$ such that

- The structure $(TX, \sqsubseteq_X, \perp_X)$ is an ωCPPO .
- The bind operator is continuous in both arguments. That is,

$$\begin{aligned} \left(\bigsqcup_{n < \omega} u_n \right) \gg f &= \bigsqcup_{n < \omega} (u_n \gg f); \\ u \gg \left(\bigsqcup_{n < \omega} f_n \right) &= \bigsqcup_{n < \omega} (u \gg f_n). \end{aligned}$$

We say that \gg is strict in its first argument if we additionally have $\perp \gg f = \perp$ (and similarly for its second argument). We say that T carries a continuous Σ -algebra structure if T has an ωCPPO order such that TX is a continuous Σ -algebra with respect to the order \sqsubseteq_X , for any set X .

Most of the time we will work with a fixed set X . As a consequence, we will omit subscripts, just writing \sqsubseteq in place of \sqsubseteq_X . Similarly, for an operation σ in Σ , we will write σ^T in place of σ^{TX} (the interpretation of σ as an operation on TX).

In the last definition we are essentially regarding the bind operator as a function from $TX \times (X \rightarrow TY)$ to TY , continuous in both arguments. This makes sense since $TX \times (X \rightarrow TY)$ is an ωCPPO (since we can regard any set X as an ωCPPO with the identity order, the set $X \rightarrow TY$ coincides with the set of continuous functions from X to TX), so that $TX \times (X \rightarrow TY)$ is indeed an ωCPPO , being the product of two ωCPPO s). Because \gg is continuous in both its arguments, we have $(\bigsqcup_n u_n) \gg (\bigsqcup_n f_n) = \bigsqcup_n (u_n \gg f_n)$.

The bind operation will be useful when giving an operational semantics to the sequential (monadic) composition of programs. As a consequence, although we did not explicitly

require the bind operator to be strict (especially in its first argument), such condition will be often desired (especially when giving semantics to call-by-value languages). It is easy to check that all bind operations defined in Example 1 are strict in their first argument.

Remark 2. The above continuity condition is a special case (since we are in \mathbb{SET}) of the more general notion of ωCPPO -enrichment. Recall that a category \mathbb{C} is ωCPPO -enriched if each hom-set $\text{Hom}_{\mathbb{C}}(X, Y)$ carries a partial order \sqsubseteq with an ωCPPO structure, and arrows' composition is continuous regarded as a function from $\text{Hom}_{\mathbb{C}}(X, Y) \times \text{Hom}_{\mathbb{C}}(Y, Z)$ to $\text{Hom}_{\mathbb{C}}(X, Z)$. A monad T on \mathbb{C} is order-enriched if the Kleisli category $\mathcal{Kl}(T)$ is ωCPPO -enriched. That is, $\text{Hom}_{\mathbb{C}}(X, TY)$ carries an ωCPPO -structure such that function composition is continuous and Kleisli star is locally continuous. Concretely, we require the following equations to hold (for arrows f, g of the appropriate type) (cf. [16]):

$$\begin{aligned} \left(\bigsqcup_{n < \omega} f_n \right) \circ g &= \bigsqcup_{n < \omega} (f_n \circ g); \\ f^\dagger \circ \bigsqcup_{n < \omega} g_n &= \bigsqcup_{n < \omega} (f^\dagger \circ g_n); \\ \left(\bigsqcup_{n < \omega} f_n \right)^\dagger &= \bigsqcup_{n < \omega} f_n^\dagger. \end{aligned}$$

Example 3. Example 1 shows that all monads in Example 2 have an ωCPPO order. Moreover, it is just a matter of simple calculations to show that all bind operators are continuous in both arguments.

IV. A COMPUTATIONAL CALCULUS AND ITS OPERATIONAL SEMANTICS

In this section we define a computational λ -calculus. Following [30], [25], [28], we syntactically distinguish between values and computations. We fix a signature Σ of operation symbols (the sources of side-effects), and a monad T carrying a continuous Σ -algebra structure (which describes the nature of the wanted effectful computations generated by the operations in Σ).

Definition 4. Given a signature Σ , the sets Λ_Σ and \mathcal{V}_Σ of terms and values are defined by the following grammars:

$$\begin{aligned} M, N &::= \text{return } V \mid VW \mid M \text{ to } x.N \mid \sigma(M, \dots, M); \\ V, W &::= x \mid \lambda x.M. \end{aligned}$$

where x ranges over a fixed countably infinite set X of variables and σ ranges over Σ .

The term $(M \text{ to } x.N)$ captures monadic binding (which is usually expressed using a “let-in” notation). A calculus with an explicit separation between terms and values has the advantage to make proofs simpler, without sacrificing expressiveness. For instance, we can encode terms' application MN as $(M \text{ to } x.(N \text{ to } y.xy))$ and vice versa $(M \text{ to } x.N)$ as $(\lambda x.N)M$.

Example 4. We can model several calculi combining the signatures from Example 1.

- For a given set E of exceptions, we can define a probabilistic λ -calculus with exceptions as Λ_Σ , for a signature $\Sigma = \{\oplus_p, \text{raise}_e \mid p \in [0, 1], e \in E\}$. In particular, we will have terms of the form $M \oplus_p N$ and raise_e . Replacing the probabilistic choice operator \oplus_p with its nondeterministic counterpart \oplus we obtain a nondeterministic calculus with exceptions.
- We can define a nondeterministic calculus with global (boolean) states as Λ_Σ , for a signature $\Sigma = \{\oplus, \text{write}_b, \text{read} \mid b \in \{\text{true}, \text{false}\}\}$. In particular, we will have terms of the form $M \oplus N$, $\text{write } b.M$, and $\text{read}(M, N)$. The intuitive meaning of $\text{write } b.M$ is to store b and then continue as M , whereas the intuitive meaning of $\text{read}(M, N)$ is to read the value in the store: if the latter is the boolean true then continue as M , otherwise as N . A formal semantics for these two functions is given in Example 1.

In what follows, we work with a fixed arbitrary signature Σ . As a consequence, we often denote the sets of terms and values as Λ and \mathcal{V} , respectively, thus omitting subscripts. Moreover, we consider terms and values modulo α -equivalence and assume Barendregt Convention [3]. We let $FV(M)$ denote the set of free variables of the term M . A term M is closed if $FV(M) = \emptyset$. We denote finite sets of variables, terms and values using “bar notation”: for instance, we write \bar{x} and \bar{V} for a finite set of variables and values, respectively. For a finite set \bar{x} of variables define

$$\begin{aligned} \Lambda(\bar{x}) &= \{M \mid FV(M) \subseteq \bar{x}\}; \\ \mathcal{V}(\bar{x}) &= \{V \mid FV(V) \subseteq \bar{x}\}; \end{aligned}$$

to be the sets of terms and values with free variables in \bar{x} , respectively. The set of closed terms and values are then defined as $\Lambda(\emptyset)$ and $\mathcal{V}(\emptyset)$, and denoted as Λ_0 and \mathcal{V}_0 , respectively. The term obtained substituting each occurrence of x by V in the term M , denoted $M[x := V]$ can be defined as usual [9].

Big-step semantics associates to each closed term M an element $\llbracket M \rrbracket$ in $T\mathcal{V}_0$. Such a semantics is defined by means of an approximation relation \Downarrow_n , indexed by a natural number n , whose definition is given in Figure 3. Judgments are of the form $M \Downarrow_n X$, where $M \in \Lambda_0$, $X \in T\mathcal{V}_0$ and $n \geq 0$. Intuitively, a judgment $M \Downarrow_n X$ states that X is the n -th approximation of the computation obtained by call-by-value evaluating M . (By the way, all the results in this paper would remain valid also if evaluating terms in call-by-name order, which is however less natural in presence of effects.)

It is easy to see that the set $\{X \in T\mathcal{V}_0 \mid M \Downarrow_n X\}$ of finite approximants of M forms an ω -chain. As a consequence, we can define the evaluation $\llbracket M \rrbracket$ of a term M as

$$\llbracket M \rrbracket = \bigsqcup_{M \Downarrow_n X} X.$$

This allows us to explicitly capture non-termination (which is usually defined coinductively). For instance, it is easy to show that for the purely (i.e. having no side-effects) divergent program Ω , we have $\llbracket \Omega \rrbracket = \perp$. This style of operational semantics

$$\boxed{\begin{array}{c} \frac{}{M \Downarrow_0 \perp} \text{(bot)} \quad \frac{}{\text{return } V \Downarrow_{n+1} \eta(V)} \text{(ret)} \\ \\ \frac{M \Downarrow_n X \quad N[x := V] \Downarrow_n Y_V}{M \text{ to } x.N \Downarrow_{n+1} X \gg=(V \mapsto Y_V)} \text{(seq)} \\ \\ \frac{M[x := V] \Downarrow_n X}{(\lambda x.M)V \Downarrow_{n+1} X} \text{(app)} \\ \\ \frac{M_1 \Downarrow_n X_1 \quad \dots \quad M_k \Downarrow_n X_k}{\sigma(M_1, \dots, M_k) \Downarrow_{n+1} \sigma^T(X_1, \dots, X_k)} \text{(op)} \end{array}}$$

Fig. 3. Big-step Semantics.

[11], [10] is precisely the reason we require the monad T to carry an ω CPPO structure. Modelling divergence explicitly turned out to be fundamental in e.g. probabilistic calculi [11].

Since both $\gg=$ and σ^T are continuous, we can characterise operational semantics equationally.

Lemma 1. *The following equations hold:*

$$\begin{aligned} \llbracket \text{return } V \rrbracket &= \eta(V); \\ \llbracket (\lambda x.M)V \rrbracket &= \llbracket M[x := V] \rrbracket; \\ \llbracket M \text{ to } x.N \rrbracket &= \llbracket M \rrbracket \gg=(V \mapsto \llbracket N[x := V] \rrbracket); \\ \llbracket \sigma(M_1, \dots, M_n) \rrbracket &= \sigma^T(\llbracket M_1 \rrbracket, \dots, \llbracket M_n \rrbracket). \end{aligned}$$

V. ON RELATIONAL REASONING

A. Relators

In this section we introduce the concept of relator [39], [27], which is an abstraction meant to capture the possible ways a relation on a set X can be turned into a relation on TX . Recall that for an endofunctor $F : \mathbb{C} \rightarrow \mathbb{C}$, an F -coalgebra [36] consists of an object X of \mathbb{C} together with a morphism $\gamma_X : X \rightarrow FX$. As usual, we are just concerned with the case in which \mathbb{C} is SET .

Definition 5. *Let F be an endofunctor on SET , and X, Y be sets. A relator Γ for F is a map that associates to each relation $\mathcal{R} \subseteq X \times Y$ a relation $\Gamma\mathcal{R} \subseteq FX \times FY$ such that*

$$=_{FX} \subseteq \Gamma(=_X) \quad (\text{Rel-1})$$

$$\Gamma\mathcal{S} \circ \Gamma\mathcal{R} \subseteq \Gamma(\mathcal{S} \circ \mathcal{R}) \quad (\text{Rel-2})$$

$$\Gamma((f \times g)^{-1}\mathcal{R}) = (Ff \times Fg)^{-1}\Gamma\mathcal{R} \quad (\text{Rel-3})$$

$$\mathcal{R} \subseteq \mathcal{S} \implies \Gamma\mathcal{R} \subseteq \Gamma\mathcal{S} \quad (\text{Rel-4})$$

where for $f : Z \rightarrow X$, $g : W \rightarrow Y$ we have $(f \times g)^{-1}\mathcal{R} = \{(z, w) \mid f(z) \mathcal{R} g(w)\}$, and $=_X$ denotes the identity relation on X . A relator Γ is conservative if $\Gamma(\mathcal{R}^c) = (\Gamma\mathcal{R})^c$, where \mathcal{R}^c denotes the converse of \mathcal{R} .

Example 5. *For each of the monads introduced in previous sections, we give some examples of relators. We use the notation Γ for a relator aimed to capture the structure of a simulation relation, and Δ for a relator aimed to capture the structure of a bisimulation relation. This distinction is*

not formal, and only makes sense in the context of concrete examples: its purpose is to stress that from formal view point, both concrete notions of similarity and bisimilarity are modeled as forms of Γ -similarity (for a suitable relator Γ). Let $\mathcal{R} \subseteq X \times Y$:

- For the partiality monad $TX = X_\perp$ define the relators $\Gamma_\perp, \Delta_\perp$ by

$$\begin{aligned} u \Gamma_\perp \mathcal{R} v \text{ iff } u = \text{in}_l(x) &\implies v = \text{in}_l(y) \wedge x \mathcal{R} y; \\ u \Delta_\perp \mathcal{R} v \text{ iff } u = \text{in}_l(x) &\implies v = \text{in}_l(y) \wedge x \mathcal{R} y, \\ v = \text{in}_l(y) &\implies u = \text{in}_l(x) \wedge x \mathcal{R} y. \end{aligned}$$

Note that $u = \text{in}_l(x)$ means, in particular, $u \neq \text{in}_r(\perp)$. Thus, for instance, u and v are $\Gamma_\perp \mathcal{R}$ related if whenever u converges, so does v and the values to which u, v converge are \mathcal{R} -related. The relator Δ_\perp is converse.

- For the nondeterministic powerset monad \mathcal{P} define relators $\Gamma_{\mathcal{P}}$ and $\Delta_{\mathcal{P}}$ by

$$\begin{aligned} u \Gamma_{\mathcal{P}} \mathcal{R} v \text{ iff } \forall x \in u. \exists y \in v. x \mathcal{R} y; \\ u \Delta_{\mathcal{P}} \mathcal{R} v \text{ iff } \forall x \in u. \exists y \in v. x \mathcal{R} y, \\ \forall y \in v. \exists x \in u. x \mathcal{R} y. \end{aligned}$$

The relator $\Delta_{\mathcal{P}}$ is converse.

- For the probabilistic subdistributions monad \mathcal{D} define relators $\Gamma_{\mathcal{D}}$ and $\Delta_{\mathcal{D}}$ by

$$\begin{aligned} \mu \Gamma_{\mathcal{D}} \mathcal{R} \nu \text{ iff } \forall U \subseteq X. \mu(U) \leq \nu(\mathcal{R}(U)); \\ \mu \Delta_{\mathcal{D}} \mathcal{R} \nu \text{ iff } \mu \Gamma_{\mathcal{D}} \mathcal{R} \nu \wedge \nu \Gamma_{\mathcal{D}} \mathcal{R}^c \mu; \end{aligned}$$

where $\mathcal{R}(U) = \{y \in Y \mid \exists x \in U. x \mathcal{R} y\}$ and $\mu(U) = \sum_{x \in U} \mu(x)$. The relator $\Delta_{\mathcal{D}}$ is converse.

- For the exception monad $TX = X + E$ define the relators $\Gamma_{\mathcal{E}}$ and $\Delta_{\mathcal{E}}$ by (letters e, e' range over E)

$$\begin{aligned} u \Gamma_{\mathcal{E}} \mathcal{R} v \text{ iff } u = \text{in}_r(e) &\implies v = \text{in}_r(e') \wedge e = e', \\ u = \text{in}_l(x) &\implies v = \text{in}_l(y) \wedge x \mathcal{R} y; \\ u \Delta_{\mathcal{E}} \mathcal{R} v \text{ iff } u \Gamma_{\mathcal{E}} \mathcal{R} v, \\ v = \text{in}_r(e') &\implies u = \text{in}_r(e) \wedge e = e', \\ v = \text{in}_l(y) &\implies u = \text{in}_l(x) \wedge x \mathcal{R} y. \end{aligned}$$

The relator $\Delta_{\mathcal{E}}$ is converse.

- For the partiality and exception monad (i.e. the exception monad with divergence) $TX = (X + E)_\perp$ we can define relators simply composing relators for the partiality monad with relators for the exceptions monads (see Lemma 2). Notably, define $\Gamma_{\mathcal{E}_\perp}$ as $\Gamma_\perp \circ \Gamma_{\mathcal{E}}$ and $\Delta_{\mathcal{E}_\perp}$ as $\Delta_\perp \circ \Delta_{\mathcal{E}}$. The relator $\Delta_{\mathcal{E}_\perp}$ is converse.

- For the state monad $TX = (X \times S)^S$ define the relator Δ_S by

$$\begin{aligned} f \Delta_S \mathcal{R} g \text{ iff } \forall s \in S. s_1 = s_2 \text{ and } x_1 \mathcal{R} x_2, \\ \text{where } (x_1, s_1) = f(s) \text{ and } (x_2, s_2) = g(s). \end{aligned}$$

The relator Δ_S is converse.

Checking that the above are indeed relators is a tedious but easy exercise. It is useful to know that the collection of relators is closed under certain operations (see [27] for proofs).

Lemma 2 (Algebra of Relators). *Let F, G be endofunctors on \mathbf{SET} . Then:*

1. The collection of relators for F is closed under arbitrary intersection and converse, where intersection of relators is defined pointwise, and the converse Γ^c of a relator Γ for F is defined by $\Gamma^c(\mathcal{R}) = (\Gamma \mathcal{R}^c)^c$. Moreover, for a relator Γ for F , $\Gamma \cap \Gamma^c$ is the greatest converse relator contained in Γ .
2. Let Γ, Γ' be relators for F, G , respectively. Then $\Gamma' \circ \Gamma$ is a relator for $G \circ F$. Moreover, if both Γ and Γ' are converse, then so is $\Gamma' \circ \Gamma$.

We can now give a general notion of (bi)simulation.

B. Bisimulation, in the Abstract

A relator Γ for a monad T expresses the observable part of the side-effects encoded by T . Its abstract nature allows to give abstract definitions of simulation and bisimulation parametric in the notion of observation given by Γ .

Definition 6. Let $\gamma_X : X \rightarrow FX, \gamma_Y : Y \rightarrow FY$ be F -coalgebras.

1. A Γ -simulation is a relation $\mathcal{R} \subseteq X \times Y$ such that

$$x \mathcal{R} y \implies \gamma_X(x) \Gamma \mathcal{R} \gamma_Y(y).$$

2. Γ -similarity $\lesssim_{X,Y}^\Gamma$ is the largest Γ -simulation.

Example 6. It is immediate to see that the corresponding notions of Γ -similarity for the (bi)simulation relators of Example 5 coincide with widely used notions of (bi)similarity.

As usual, the notion of similarity can be characterised coinductively as the greatest fixed point of a suitable functional.

Definition 7. Let $\gamma_X : X \rightarrow FX, \gamma_Y : Y \rightarrow FY$ be F -coalgebras. Define the functional $\mathcal{F}_{X,Y}^\Gamma : 2^{X \times Y} \rightarrow 2^{X \times Y}$ by

$$\mathcal{F}_{X,Y}^\Gamma(\mathcal{R}) = (\gamma_X \times \gamma_Y)^{-1}(\Gamma \mathcal{R}).$$

When clear from the context, we will write \mathcal{F}_Γ and \lesssim_Γ in place of $\mathcal{F}_{X,Y}^\Gamma$ and $\lesssim_{X,Y}^\Gamma$.

Lemma 3. *The following hold:*

1. The functional \mathcal{F}_Γ is monotone, and thus has a greatest fixed point $\nu \mathcal{F}_\Gamma$.
2. A relation \mathcal{R} is a Γ -simulation iff it is a post fixed-point of \mathcal{F}_Γ . Therefore, Γ -similarity coincides with $\nu \mathcal{F}_\Gamma$.

It is easy to see that (Rel-1) implies that the identity relation is a Γ -simulation, whereas property (Rel-2) makes Γ -similarity transitive. Moreover, for a converse relator Γ we have that if \mathcal{R} is a Γ -simulation, then so is \mathcal{R}^c . As a consequence, we have the following.

Proposition 1. *Let $\gamma_X : X \rightarrow FX$ be an F -coalgebra.*

1. Γ -similarity is a preorder.
2. If Γ is converse, then Γ -similarity is an equivalence relation.

Since T is a monad we consider relators that properly interact with the monadic structure of T , which are also known as *lax extensions* for T [4].

Definition 8. Let T be a monad, X, X', Y, Y' be sets, $f : X \rightarrow TX', g : Y \rightarrow TY'$ be functions, and $\mathcal{R} \subseteq X \times Y, \mathcal{S} \subseteq X' \times Y'$ be relations. We say that Γ is a relator for T if it is a relator for T regarded as a functor, and

- $x \mathcal{R} y \implies \eta_X(x) \Gamma \mathcal{R} \eta_Y(y)$;
- $u \Gamma \mathcal{R} v \implies (u \gg f) \Gamma \mathcal{S} (v \gg g)$, whenever $x \mathcal{R} y \implies f(x) \Gamma \mathcal{S} g(y)$.

Remark 3. Definition 8 can be more compactly expressed using Kleisli star, thus requiring that

$$\begin{aligned} \mathcal{R} &\subseteq (\eta_X \times \eta_Y)^{-1}(\Gamma \mathcal{S}) && \text{(Lax-Unit)} \\ \mathcal{R} &\subseteq (f \times g)^{-1}(\Gamma \mathcal{S}) \implies \Gamma \mathcal{R} \subseteq (f^\dagger \times g^\dagger)^{-1}(\Gamma \mathcal{S}) && \text{(Lax-Bind)} \end{aligned}$$

or diagrammatically

$$\begin{array}{ccc} X & \xrightarrow{\mathcal{R}} & Y \\ \eta_X \downarrow & & \downarrow \eta_Y \\ TX & \xrightarrow{\Gamma \mathcal{R}} & TY \end{array} \quad \begin{array}{ccc} X & \xrightarrow{\mathcal{R}} & Y \\ f \downarrow & & \downarrow g \\ TX' & \xrightarrow{\Gamma \mathcal{S}} & TY' \end{array} \implies \begin{array}{ccc} TX & \xrightarrow{\Gamma \mathcal{R}} & TY \\ f^\dagger \downarrow & & \downarrow g^\dagger \\ TX & \xrightarrow{\Gamma \mathcal{S}} & TY \end{array}$$

where we write $\mathcal{R} : X \rightsquigarrow Y$ for $\mathcal{R} \subseteq X \times Y$.

Example 7. All relators of the form Γ_T in Example 5 are relators for T . Proving that is quite standard, with the exception of the probabilistic case where the proof essentially relies on the Max Flow Min Cut Theorem [38].

Definition 9. Let T come with an ω CPPO order \sqsubseteq . We say that $\Gamma \mathcal{R}$ is inductive if for any ω -chain $(u_n)_{n < \omega}$ in TX , we have:

$$\begin{aligned} \perp \Gamma \mathcal{R} u; & && (\omega\text{-comp } 1) \\ (\forall n. u_n \Gamma \mathcal{R} v) \implies \bigsqcup_n u_n \Gamma \mathcal{R} v. & && (\omega\text{-comp } 2) \end{aligned}$$

We say that Γ respects Σ if

$$(\forall k. u_k \Gamma \mathcal{R} v_k) \implies \sigma(u_1, \dots, u_n) \Gamma \mathcal{R} \sigma(v_1, \dots, v_n) \quad (\Sigma\text{-comp})$$

for any $\sigma \in \Sigma$, where $k \in \{1, \dots, \alpha(\sigma)\}$.

Remark 4. For a monad T carrying a continuous Σ -algebra structure and a function $f : X \rightarrow TY$, we required $f^\dagger : TX \rightarrow TY$ to be continuous, TX being an ω CPPO. Since TX is also a Σ -algebra, it seems natural to require f^\dagger to be also a Σ -algebra homomorphism. In fact, such requirement implies condition $(\Sigma\text{-comp})$ and has the advantage of being more general than the latter, not depending from the specific relator considered (see [9] for more details). To the ends of this paper, condition $(\Sigma\text{-comp})$ is sufficient and thus we will use that throughout.

Following Abramsky [1] we introduce Applicative Transition System (ATSs) over a monad (taking into account effectful

computations) and define the notion of applicative simulation. Let T be a monad.

Definition 10. An applicative transition system (over T) consists of the following:

- A state space made of a pair of sets (X, Y) modelling closed terms and values, respectively.
- An evaluation function $\varepsilon : X \rightarrow TY$.
- An application function $\cdot : Y \rightarrow Y \rightarrow X$.

The notion of ATS distinguishes between terms and values. As a consequence, we often deal with pairs of relations $(\mathcal{R}_X, \mathcal{R}_Y)$, where $\mathcal{R}_X, \mathcal{R}_Y$ are relations over X and Y , respectively. We refer to such pairs as XY -relations. XY -relations belongs to $2^{X \times X} \times 2^{Y \times Y}$. The latter, being the product of complete lattices, is itself a complete lattice.

Definition 11. Let Γ be a relator for T . An applicative Γ -simulation is an XY -relation $\mathcal{R} = (\mathcal{R}_X, \mathcal{R}_Y)$ such that:

$$x \mathcal{R}_X x' \implies \varepsilon(x) \Gamma \mathcal{R}_Y \varepsilon(x') \quad (\text{Sim-1})$$

$$y \mathcal{R}_Y y' \implies \forall w \in Y. y \cdot w \mathcal{R}_X y' \cdot w \quad (\text{Sim-2})$$

The above definition induces an operator \mathcal{B}_Γ on $2^{X \times X} \times 2^{Y \times Y}$ defined for $\mathcal{R} = (\mathcal{R}_X, \mathcal{R}_Y)$ as $(\mathcal{B}_\Gamma(\mathcal{R}_X), \mathcal{B}_\Gamma(\mathcal{R}_Y))$, where

$$\mathcal{B}_\Gamma(\mathcal{R}_X) = \{(x, x') \mid \varepsilon(x) \Gamma \mathcal{R}_Y \varepsilon(x')\};$$

$$\mathcal{B}_\Gamma(\mathcal{R}_Y) = \{(y, y') \mid \forall w \in Y. y \cdot w \mathcal{R}_X y' \cdot w\}.$$

It is easy to prove that since Γ is monotone, then so is \mathcal{B}_Γ . As a consequence, we can define applicative Γ -similarity as the greatest fixed point $\nu \mathcal{B}_\Gamma$ of \mathcal{B}_Γ .

Proposition 2. The following hold:

1. Applicative Γ -similarity $\nu \mathcal{B}_\Gamma$ is a preorder.
2. If Γ is converse, then $\nu \mathcal{B}_\Gamma$ is an equivalence relation.

VI. CONTEXTUAL PREORDER AND APPLICATIVE SIMILARITY

In the previous section, the axioms needed to generalise applicative bisimilarity to our setting have been given. What remains to be done is to appropriately instantiate all this to Λ_Σ . We introduce the notions of contextual preorder and applicative similarity (which will be then extended to contextual equivalence and applicative bisimilarity). From now we assume to have a monad T carrying a continuous Σ -algebra structure. Moreover, we assume any relator for T to be inductive and to respect Σ . It is convenient to work with generalisations of relations on closed terms (resp. values) called λ -term relations.

Definition 12. An open relation over terms is a set \mathcal{R}_Λ of triples (\bar{x}, M, N) where $M, N \in \Lambda(\bar{x})$. Similarly, an open relation over values is a set $\mathcal{R}_\mathcal{V}$ of triples (\bar{x}, V, W) where $V, W \in \mathcal{V}(\bar{x})$. A λ -term relation is a pair $\mathcal{R} = (\mathcal{R}_\Lambda, \mathcal{R}_\mathcal{V})$ made of an open relation \mathcal{R}_Λ over terms and an open relation $\mathcal{R}_\mathcal{V}$ over values. A closed λ -term relation is a pair $\mathcal{R} = (\mathcal{R}_\Lambda, \mathcal{R}_\mathcal{V})$ where $\mathcal{R}_\Lambda \subseteq \Lambda_0 \times \Lambda_0$ and similarly for $\mathcal{R}_\mathcal{V}$.

$\forall x \in \bar{x}. \bar{x} \vdash x \mathcal{R}_V x$	(Comp1)
$\forall x \notin \bar{x}. \bar{x} \cup \{x\} \vdash M \mathcal{R}_\Lambda N \implies \bar{x} \vdash \lambda x.M \mathcal{R}_V \lambda x.N$	(Comp2)
$\bar{x} \vdash V \mathcal{R}_V W \implies \bar{x} \vdash \text{return } V \mathcal{R}_\Lambda \text{return } W$	(Comp3)
$\bar{x} \vdash V \mathcal{R}_V V' \wedge \bar{x} \vdash W \mathcal{R}_V W' \implies \bar{x} \vdash VW \mathcal{R}_\Lambda V'W'$	(Comp4)
$\forall x \notin \bar{x}. \bar{x} \vdash M \mathcal{R}_\Lambda M' \wedge \bar{x} \cup \{x\} \vdash N \mathcal{R}_\Lambda N' \implies \bar{x} \vdash (M \text{ to } x.N) \mathcal{R}_\Lambda (M' \text{ to } x.N')$	(Comp5)
$\bar{x} \vdash M_1 \mathcal{R}_\Lambda N_1 \wedge \dots \wedge \bar{x} \vdash M_{\alpha(\sigma)} \mathcal{R}_\Lambda N_{\alpha(\sigma)} \implies \bar{x} \vdash \sigma(M_1, \dots, M_{\alpha(\sigma)}) \mathcal{R}_\Lambda \sigma(N_1, \dots, N_{\alpha(\sigma)})$	(Comp6)

Fig. 4. Compatibility Clauses.

Remark 5. Formally, we can see an open relation over terms (and similarly over values) as an element of the cartesian product $\prod_{\bar{x}} 2^{\Lambda(\bar{x}) \times \Lambda(\bar{x})}$. That is, an open relation is a function that associates to each finite set \bar{x} of variables a (binary) relation between open terms in $\Lambda(\bar{x})$. Since, $2^{\Lambda(\bar{x}) \times \Lambda(\bar{x})}$ is a complete lattice, for any finite set of variables \bar{x} , then so is $\prod_{\bar{x}} 2^{\Lambda(\bar{x}) \times \Lambda(\bar{x})}$. That is, the set of open relations over terms (and over values) forms a complete lattice (the order is given pointwise). As a consequence, the set of λ -term relations is a complete lattice as well. These algebraic properties allow us to define open relations both inductively and coinductively, and, in particular, to extend notions and results developed in the relational calculus [25], [24], [17], [26].

We will use infix notation and write $\bar{x} \vdash M \mathcal{R}_\Lambda N$ to indicate that $(\bar{x}, M, N) \in \mathcal{R}_\Lambda$. The same convention applies to values and open relations over values. For a λ -term relation $\mathcal{R} = (\mathcal{R}_\Lambda, \mathcal{R}_V)$, we often write $\bar{x} \vdash M \mathcal{R} N$ (i.e. $(\bar{x}, M, N) \in \mathcal{R}$) for $\bar{x} \vdash M \mathcal{R}_\Lambda N$ (i.e. $(\bar{x}, M, N) \in \mathcal{R}_\Lambda$). The same convention holds for values and \mathcal{R}_V . Finally, we will use the notations $\emptyset \vdash M \mathcal{R} N$ and $M \mathcal{R} N$ interchangeably (and similarly for values).

There is a canonical way to extend a closed relation to an open one.

Definition 13. Define the open extension operator mapping a closed relation over terms \mathcal{R} to the open relation \mathcal{R}° (over terms) as follows: $(\bar{x}, M, N) \in \mathcal{R}^\circ$ iff $M, N \in \Lambda(\bar{x})$, and for all \bar{V} , $M[\bar{x} := \bar{V}] \mathcal{R} N[\bar{x} := \bar{V}]$ holds.

The notion of open extension for a closed relation over values can be defined in a similar way (using the appropriate notion of substitution).

The notion of reflexivity, symmetry and transitivity straightforwardly extends to open λ -term relation (see e.g. [33]).

Definition 14. Let $\mathcal{R} = (\mathcal{R}_\Lambda, \mathcal{R}_V)$ be a λ -term relation. We say that \mathcal{R} is compatible if the clauses in Figure 4 hold. We say that \mathcal{R} is a precongruence if it is a compatible preorder. We say that \mathcal{R} is a congruence if it is a compatible equivalence.

It is useful to characterise compatible relations via the notion of compatible refinement.

Definition 15. Let $\mathcal{R} = (\mathcal{R}_\Lambda, \mathcal{R}_V)$ be a λ -term relation. De-

fine the compatible refinement $\widehat{\mathcal{R}}$ of \mathcal{R} as the pair $(\widehat{\mathcal{R}}_\Lambda, \widehat{\mathcal{R}}_V)$, where \mathcal{R}_Λ and \mathcal{R}_V are inductively defined by rules in Figure 5. It is easy to see that a λ -term relation \mathcal{R} is compatible iff $\widehat{\mathcal{R}} \subseteq \mathcal{R}$ holds.

$\frac{}{\bar{x} \vdash x \widehat{\mathcal{R}}_V x} \quad x \in \bar{x}$	$\frac{\bar{x} \cup \{x\} \vdash M \mathcal{R}_\Lambda N}{\bar{x} \vdash \lambda x.M \widehat{\mathcal{R}}_V \lambda x.N} \quad x \notin \bar{x}$
$\bar{x} \vdash V \mathcal{R}_V W$	
$\bar{x} \vdash \text{return } V \widehat{\mathcal{R}}_\Lambda \text{return } W$	
$\bar{x} \vdash V \mathcal{R}_V V'$	$\bar{x} \vdash W \mathcal{R}_V W'$
$\bar{x} \vdash VW \widehat{\mathcal{R}}_\Lambda V'W'$	
$\bar{x} \vdash M \mathcal{R}_\Lambda M'$	$\bar{x} \cup \{x\} \vdash N \mathcal{R}_\Lambda N'$
$\bar{x} \vdash M \text{ to } x.N \widehat{\mathcal{R}}_\Lambda M' \text{ to } x.N'$	
$\bar{x} \vdash M_1 \mathcal{R}_\Lambda N_1$	$\dots \quad \bar{x} \vdash M_n \mathcal{R}_\Lambda N_n$
$\bar{x} \vdash \sigma(M_1, \dots, M_n) \widehat{\mathcal{R}}_\Lambda \sigma(N_1, \dots, N_n)$	

Fig. 5. Compatible Refinement Rules.

The above notion of precongruence can be justified by observing that when a relation \mathcal{R} is a preorder, being a precongruence does exactly mean to be closed under the term constructors of the language. That could be formally expressed by saying that \mathcal{R} is a precongruence if and only if $\bar{x} \vdash M \mathcal{R} N$ implies $\bar{x} \vdash C[M] \mathcal{R} C[N]$, for any term context $C[\cdot]$. Defining term contexts requires some care. In particular, when dealing with the contextual preorder it is not possible to reason modulo α -conversion, thus making definition syntactically involved (see [25], [24], [33] for details). As remarked in [33], it is possible to avoid those difficulties by giving a coinductive characterisation of the contextual preorder in the style of [25], [17]. Essentially, the contextual preorder (and, similarly the contextual equivalence) is defined as the largest compatible and preadequate (see Definition 16) λ -term relation. It is then easy to provide a more syntactic definition of contextual preorder and to prove that the two given definitions are equivalent [17], [25], [33].

The notion of adequacy defines the available observation on values. Being in an untyped setting, it is customary not to observe them.

Definition 16. Let \mathcal{U} denote $\mathcal{V}_0 \times \mathcal{V}_0$ seen as a closed relation, i.e. the trivial relation relating all values. We say that a relation \mathcal{R} on terms is preadequate if

$$\emptyset \vdash M \mathcal{R} N \implies \llbracket M \rrbracket \Gamma \llbracket N \rrbracket;$$

where $M, N \in \Lambda_0$. That is, a relation \mathcal{R} on terms is preadequate if whenever \mathcal{R} relates two closed terms, evaluating these programs produces the same side-effects. A λ -term relation $\mathcal{R} = (\mathcal{R}_\Lambda, \mathcal{R}_V)$ is preadequate iff \mathcal{R}_Λ is.

Example 8. It is easy to check that the above notion of adequacy (together with the relators in Example 5) captures standard notions of adequacy used for untyped λ -calculi.

VII. HOWE'S METHOD AND ITS SOUNDNESS

- Consider a calculus without operation symbols and with operational semantics over $(\mathcal{V}_0)_\perp$. A relation is preadequate if whenever $\emptyset \vdash M \mathcal{R} N$, then if M converges, then so does N .
- Consider a nondeterministic calculus with operational semantics over \mathcal{PV}_0 . A relation is preadequate if whenever $\emptyset \vdash M \mathcal{R} N$, then if there exists a value V to which M may converge (i.e. $V \in \llbracket M \rrbracket$), then there exists a value W to which N may converge (i.e. $W \in \llbracket N \rrbracket$).
- Consider a probabilistic calculus with operational semantics over \mathcal{DV}_0 . A relation is preadequate if whenever $\emptyset \vdash M \mathcal{R} N$, then the probability of convergence of M is smaller or equal than the probability of convergence of N .

Following [25], we shall define the Γ -contextual preorder as the largest λ -term relation that is both compatible and preadequate.

Definition 17. Let \mathbb{CA} be the set of relations on terms that are both compatible and preadequate. Then define \leq_Γ as $\bigcup \mathbb{CA}$.

Proposition 3. The Γ -contextual preorder \leq_Γ is a compatible and preadequate preorder.

Finally, we can define the notion of an applicative simulation observing how the collection of closed terms and values, together with the operational semantics defined in previous section, carry an ATS structure.

Definition 18. A closed relation $\mathcal{R} = (\mathcal{R}_\Lambda, \mathcal{R}_\mathcal{V})$ respects values if for all closed values V, W , $V \mathcal{R}_\mathcal{V} W$ implies $VU \mathcal{R}_\Lambda WU$, for any closed value U .

We can now define an ATS over closed λ -terms taking the pair $(\Lambda_0, \mathcal{V}_0)$ as state space and the map $[\cdot] : \Lambda_0 \rightarrow T\mathcal{V}_0$ as evaluation function. Instantiating the general definition of applicative Γ -simulation we obtain:

Definition 19. Let Γ be a relator for the monad T . A closed relation $\mathcal{R} = (\mathcal{R}_\Lambda, \mathcal{R}_\mathcal{V})$ is an applicative Γ -simulation if:

- $M \mathcal{R}_\Lambda N \implies \llbracket M \rrbracket \Gamma \mathcal{R}_\mathcal{V} \llbracket N \rrbracket$;
- \mathcal{R} respects values.

We can then define applicative Γ -similarity \lesssim_Γ as the largest applicative Γ -simulation, which we know to be a preorder by Proposition 2. Most of the time the relator Γ will be fixed; in those cases we will often write \lesssim in place of \lesssim_Γ .

Example 9. It is immediate to see that using the relators in Example 5 we recover well-known notions of simulation and bisimulation.

We want to prove that applicative similarity is a sound proof technique for contextual preorder. That is, we want to prove that $\lesssim_\Gamma \subseteq \leq_\Gamma$ holds. The relation \leq_Γ being defined as the largest preadequate compatible relation, the above inclusion is established by proving that \lesssim_Γ is a precongruence.

In this section we generalise Howe's technique to show that applicative similarity is a precongruence, thus a sound proof technique for the contextual preorder. Our generalisation shows how Howe's method crucially (but only!) depends on the structure of the monad modelling side-effects and the relators encoding their associated notion of observation.

Definition 20. Let \mathcal{R} be a closed λ -term relation. The Howe extension \mathcal{R}^H of \mathcal{R} is defined as the least relation \mathcal{S} such that $\mathcal{S} = \mathcal{R}^\circ \circ \hat{\mathcal{S}}$.

It was observed in [26] that the above equation actually defines a unique relation.

Lemma 4. Let \mathcal{R} be a closed λ -term relation. Then there is a unique relation \mathcal{S} such that $\mathcal{S} = \mathcal{R}^\circ \circ \hat{\mathcal{S}}$.

As a consequence, \mathcal{R}^H can be characterised both inductively and coinductively.

Here we give the well-known inductive characterisations of \mathcal{R}^H as the pair $(\mathcal{R}_\Lambda^H, \mathcal{R}_\mathcal{V}^H)$ inductively defined by rules in Figure 6.

The following lemma states some nice properties of Howe's lifting of preorder relations. The proof is standard and can be found in, e.g., [10].

Lemma 5. Let \mathcal{R} be a preorder. The following hold:

1. $\mathcal{R} \circ \mathcal{R}^H \subseteq \mathcal{R}^H$.
2. \mathcal{R}^H is compatible, and thus reflexive.
3. $\mathcal{R} \subseteq \mathcal{R}^H$.

We now consider the Howe extension \lesssim_Γ^H of applicative Γ -similarity. Since \lesssim_Γ is a preorder (Proposition 2), \lesssim_Γ^H is a compatible relation containing \lesssim_Γ .

Definition 21. A λ -term relation $\mathcal{R} = (\mathcal{R}_\Lambda, \mathcal{R}_\mathcal{V})$ is value-substitutive if $x \vdash M \mathcal{R}_\Lambda N$ and $\emptyset \vdash V \mathcal{R}_\mathcal{V} W$ imply $\emptyset \vdash M[x := V] \mathcal{R}_\Lambda N[x := W]$.

Lemma 6. The relation \lesssim_Γ^H is value-substitutive.

Summing up, we have defined a compatible relation \lesssim_Γ^H which is value-substitutive and contains \lesssim_Γ . As a consequence, to prove that the latter is compatible it is sufficient to prove $\lesssim_\Gamma^H \subseteq \lesssim_\Gamma$. We can proceed coinductively, showing that \lesssim_Γ^H is an applicative Γ -simulation. This is proved via the so-called Key Lemma. It is useful to spell out basic facts on the Howe extension of applicative similarity that we will extensively use. Let Γ be a relator (in the following we assume to have fixed a relator Γ , thus omitting subscripts).

Lemma 7. Let Γ be a fixed relator and write \lesssim for \lesssim_Γ . Then

1. $\lesssim \circ \lesssim^H \subseteq \lesssim^H$.
2. $(\Gamma \lesssim) \circ (\Gamma \lesssim^H) \subseteq \Gamma \lesssim^H$.

We can now state and prove the Key Lemma.

Lemma 8 (Key Lemma). Let $\lesssim^H = (\lesssim_\Lambda^H, \lesssim_\mathcal{V}^H)$ be the Howe extension of applicative similarity. If $\emptyset \vdash M \lesssim_\Lambda^H N$ and $M \Downarrow_n X$, then $X \Gamma \lesssim_\mathcal{V}^H \llbracket N \rrbracket$.

$$\begin{array}{c}
\frac{\bar{x} \vdash x \mathcal{R}_V^\circ V}{\bar{x} \vdash x \mathcal{R}_V^H V} \text{ (How1)} \qquad \frac{\bar{x} \cup \{x\} \vdash M \mathcal{R}_\Lambda^H L \quad \bar{x} \vdash \lambda x.L \mathcal{R}_V^\circ V}{\bar{x} \vdash \lambda x.M \mathcal{R}_V^H V} \text{ (How2)} \\
\frac{\bar{x} \vdash V \mathcal{R}_V^H W \quad \bar{x} \vdash \text{return } W \mathcal{R}_\Lambda^\circ N}{\bar{x} \vdash \text{return } V \mathcal{R}_\Lambda^H N} \text{ (How3)} \quad \frac{\bar{x} \vdash V \mathcal{R}_V^H V' \quad \bar{x} \vdash W \mathcal{R}_V^H W' \quad \bar{x} \vdash V'W' \mathcal{R}_\Lambda^\circ N}{\bar{x} \vdash VW \mathcal{R}_\Lambda^H N} \text{ (How4)} \\
\frac{\bar{x} \vdash M \mathcal{R}_\Lambda^H L \quad \bar{x} \cup \{x\} \vdash M' \mathcal{R}_\Lambda^H L' \quad \bar{x} \vdash L \text{ to } x.L' \mathcal{R}_\Lambda^\circ N}{\bar{x} \vdash M \text{ to } x.M' \mathcal{R}_\Lambda^H N} \text{ (How5)} \\
\frac{\bar{x} \vdash M_k \mathcal{R}_\Lambda^H N_k \ (\forall k \geq n) \quad \bar{x} \vdash \sigma(N_1, \dots, N_n) \mathcal{R}_\Lambda^\circ N}{\bar{x} \vdash \sigma(M_1, \dots, M_n) \mathcal{R}_\Lambda^H N} \text{ (How6)}
\end{array}$$

Fig. 6. Howe's Extension Rules.

Proof. The proof proceeds by induction on the derivation of the judgment $M \Downarrow_n X$. Some interesting cases follow:

Case (app). Suppose the judgment $(\lambda x.M)V \Downarrow_{n+1} X$ has been obtained from the judgment $M[x := V] \Downarrow_n X$. By hypothesis we have $\emptyset \vdash (\lambda x.M)V \lesssim_\Lambda^H N$, meaning that the latter must have been obtained as the conclusion of an instance of (How4). We thus obtain $\emptyset \vdash \lambda x.M \lesssim_V^H W$, $\emptyset \vdash V \lesssim_V^H U$ and $WU \lesssim_\Lambda N$, for values W, U . Looking at the first of these three judgments, we see that it must be the conclusion of an instance of rule (How2). Therefore, we have $\{x\} \vdash M \lesssim_\Lambda^H L$ and $\lambda x.L \lesssim_V W$. Since \lesssim^H is value-substitutive, from $\{x\} \vdash M \lesssim_\Lambda^H L$ and $\emptyset \vdash V \lesssim_V^H U$ we conclude $\emptyset \vdash M[x := V] \lesssim_\Lambda^H L[x := U]$. We can now apply the induction hypothesis on the latter and $M[x := V] \Downarrow_n X$, obtaining $X \Gamma(\lesssim_V^H) \llbracket L[x := U] \rrbracket$. Since \lesssim respects values, from $\lambda x.L \lesssim_V W$ we infer $(\lambda x.L)U \lesssim_\Lambda WU$, which gives, by very definition of applicative similarity, $\llbracket (\lambda x.L)U \rrbracket \Gamma(\lesssim_V) \llbracket WU \rrbracket$. By Lemma 1, $\llbracket (\lambda x.L)U \rrbracket = \llbracket L[x := U] \rrbracket$, and thus, $X \Gamma(\lesssim_V^H) \llbracket WU \rrbracket$, by Lemma 7. Finally, from $WU \lesssim_V N$ we obtain $\llbracket WU \rrbracket \Gamma(\lesssim_V) \llbracket N \rrbracket$, which allows us to conclude $X \Gamma(\lesssim_V) \llbracket N \rrbracket$ by Lemma 7.

Case (seq). Suppose the judgment $(M \text{ to } x.M') \Downarrow_{n+1} X \gg (V \mapsto Y_V)$ has been obtained from $M \Downarrow_n X$ and $M'[x := V] \Downarrow_n Y_V$. By hypothesis we have $\emptyset \vdash M \text{ to } x.M' \lesssim_\Lambda^H N$, which must have been obtained via an instance of rule (How5) thus giving $\emptyset \vdash M \lesssim_\Lambda^H L$, $\{x\} \vdash M' \lesssim_\Lambda^H L'$ and $\emptyset \vdash L \text{ to } x.L' \lesssim_\Lambda N$. We can apply the induction hypothesis on $M \Downarrow_n X$ and $\emptyset \vdash M \lesssim_\Lambda^H L$ obtaining $X \Gamma(\lesssim_V^H) \llbracket L \rrbracket$. We now claim to have

$$X \gg (V \mapsto Y_V) \Gamma(\lesssim_V^H) \llbracket L \rrbracket \gg (V \mapsto \llbracket L'[x := V] \rrbracket).$$

The latter is equal to $\llbracket L \text{ to } x.L' \rrbracket$, by Lemma 1. Besides, $\emptyset \vdash L \text{ to } x.L' \lesssim_\Lambda N$ entails $\llbracket L \text{ to } x.L' \rrbracket \Gamma(\lesssim_V) \llbracket N \rrbracket$: we conclude $X \gg (V \mapsto Y_V) \Gamma(\lesssim_V^H) \llbracket N \rrbracket$, by Lemma 7.

The above claim directly follows from (Lax-Bind). In fact, since $X \Gamma(\lesssim_V^H) \llbracket L \rrbracket$ holds, by (Lax-Bind) it is sufficient to prove that $V \lesssim_V^H W$ implies $Y_V \Gamma(\lesssim_V^H) \llbracket L'[x := W] \rrbracket$. Assume $V \lesssim_V^H W$, i.e. $\emptyset \vdash V \lesssim_V^H W$. The latter, together with $\{x\} \vdash M' \lesssim_\Lambda^H L'$, implies $\emptyset \vdash M'[x := V] \lesssim_\Lambda^H L'[x := W]$, since \lesssim^H is value-substitutive. We can finally

apply the inductive hypothesis on the latter and $M'[x := V] \Downarrow_n Y_V$, thus concluding the wanted thesis. \square

Corollary 1. *The relation \lesssim_Γ^H is an applicative Γ -simulation.*

Proof. Suppose $M \lesssim_\Lambda^H N$. We have to prove $\llbracket M \rrbracket \Gamma \lesssim_V^H \llbracket N \rrbracket$, i.e. $\bigsqcup_{M \Downarrow_n X} X \Gamma \lesssim_V^H \llbracket N \rrbracket$. The latter follows from (ω -comp 1) by the Key Lemma. Finally, since \lesssim^H is compatible, it clearly respects values. \square

Theorem 1. *Similarity is a precongruence. Moreover, it is sound for contextual preorder \leq_Γ .*

Proof. We already know \lesssim_Γ is a preorder. By previous corollary it follows that \lesssim_Γ coincides with \lesssim_Γ^H , so that \lesssim_Γ is also compatible, and thus a precongruence. Now for soundness. We have to prove $\lesssim_\Gamma \subseteq \leq_\Gamma$. Since \leq_Γ is defined as the largest preadequate compatible relation, it is sufficient to prove that \lesssim_Γ is preadequate (we have already showed it is compatible), which directly follows from Sim-1, since the universal relation \mathcal{U} contains \lesssim_V . \square

VIII. BISIMILARITY, TWO-SIMILARITY AND CONTEXTUAL EQUIVALENCE

In this section we extend previous definitions and results to come up with sound proof techniques for contextual *equivalence*. In particular, by observing that contextual equivalence always coincides with the intersection between the contextual preorder and its converse, Theorem 1 implies that two-way similarity (i.e. the intersection between applicative similarity and its converse) is contained in contextual equivalence. Applicative bisimilarity being finer than two-way similarity, we can also conclude the former to be a sound proof technique for contextual equivalence.

Given a relator Γ , we can extract a canonical notion of Γ -bisimulation from the one of Γ -simulation following the idea that a bisimulation is a relation \mathcal{R} such that both \mathcal{R} and \mathcal{R}^c are simulations. Recall that given a relator Γ we can define a converse operation Γ^c as $\Gamma^c(\mathcal{R}) = (\Gamma(\mathcal{R}^c))^c$. Γ^c is indeed a relator. Similarly, we have proved that the intersection of relators is again a relator.

Definition 22 (Γ -bisimulation). *Given a relator Γ , we say that a relation \mathcal{R} is a Γ -bisimulation if it is a $(\Gamma \cap \Gamma^c)$ -simulation.*

From the above definition, it directly follows that a relation \mathcal{R} is a Γ -bisimulation if and only if both \mathcal{R} and \mathcal{R}^c are Γ -simulation.

Since, by Lemma 2, $\Gamma \cap \Gamma^c$ is a relator, we can define Γ -bisimilarity \sim_Γ as $(\Gamma \cap \Gamma^c)$ -similarity. It is easy to see that \sim_Γ is an equivalence relation (note that if \mathcal{R} is a Γ -simulation, then so is Γ^c).

Definition 23. *Let Γ be a relator. Define Γ -cosimilarity \succsim_Γ as $(\succsim_\Gamma)^c$. Define Γ two-way similarity \simeq_Γ as $\succsim_\Gamma \cap \sim_\Gamma$.*

As usual, bisimilarity is finer than two-way similarity, meaning that $\sim_\Gamma \subseteq \simeq_\Gamma$. Moreover, taking Γ to be the simulation relator for the powerset monad (see Example 5), we have that \sim_Γ and \simeq_Γ do not coincide. See e.g. [25], [33].

Recall that we have defined the Γ -contextual preorder \leq_Γ as the largest relation that is both compatible and Γ -preadequate. In analogy with what we did for simulation and bisimulation we can give the following:

Definition 24. *Let Γ be a relator. Define Γ -contextual equivalence \equiv_Γ as the largest relation that is both compatible and $(\Gamma \cap \Gamma^c)$ -preadequate. That is, define \equiv_Γ as $\leq_{\Gamma \cap \Gamma^c}$.*

Although bisimilarity is finer than two-way similarity, this is not the case for contextual equivalence and the associated contextual preorders.

Proposition 4. *Let Γ be a relator. Then, $\equiv_\Gamma = \leq_\Gamma \cap \geq_\Gamma$.*

We can finally prove our soundness result.

Theorem 2 (Soundness). *Let Γ be a relator. Two-way similarity \simeq_Γ is a congruence, and thus sound for contextual equivalence \equiv_Γ . Since bisimilarity \sim_Γ is finer than \simeq_Γ , it is sound for \equiv_Γ as well.*

Proof. From Theorem 1 we know that \succsim_Γ is a precongruence and that $\succsim_\Gamma \subseteq \leq_\Gamma$. It follows \sim_Γ is a precongruence as well, and that $\sim_\Gamma \subseteq \geq_\Gamma$ holds. We can conclude \simeq_Γ is a congruence and $\simeq_\Gamma \subseteq \equiv_\Gamma$. Since $\sim_\Gamma \subseteq \simeq_\Gamma$, we also have $\sim_\Gamma \subseteq \equiv_\Gamma$. \square

Noticeably, Theorem 2 can be seen as a proof of soundness for applicative bisimilarity in any calculus Λ_Σ which respects our requirements (see Definition 8, 9), and in particular for those described in Example 5. The case of probabilistic calculi is illuminating: the apparent complexity of all proofs of congruence from the literature [10], [8] has been confined to the proof that the relator for subdistributions satisfies our axiomatics.

Theorem 2 also allows us to prove W^{raise} and Z^{raise} , our example programs from Section II, to be equivalent. This only requires checking that the map $\Gamma_{\mathcal{D}} \circ \Gamma_{\mathcal{E}}$ (see Example 5) is an inductive relator for the monad $TX = \mathcal{D}(X + E)$ (which trivially carries a continuous Σ -algebra structure) respecting operations in Σ . This turned out to be an easy exercise. Details can be found in [9].

As mentioned in the Introduction, this is certainly not the first paper about program equivalence for higher-order effectful calculi. Denotational semantics of calculi having this nature, has been studied since Moggi's seminal work [30], thus implicitly providing a notion of equivalence. All this has been given a more operational flavour starting with Plotkin and Power account on adequacy for algebraic effects [34], from which the operational semantics presented in this paper is greatly inspired. The literature also offers abstract accounts on logical relations for effectful calculi. The first of them is due to Goubault-Larrecq, Lasota and Nowak [18], which is noticeably able to deal with nondeterministic and probabilistic effects, but also with dynamic name creation, for which applicative bisimilarity is known to be unsound. Another piece of work which is related to ours is due to Johann, Simpson, and Voigtländer [20], who focused on algebraic effects and observational equivalence, and their characterisation via CIU theorems and a form of logical relation based TT -lifting. In both cases, the target language is typed. Similar in spirit to our approach (which is based on the notion of relator), the work of Katsumata and Sato [22] analyses monadic lifting of relations in the context of TT -lifting.

Although no abstract account exists on applicative coinductive techniques for calculi with algebraic effects, some work definitely exists in some specific cases. As a noticeable example, the works by Ong [32] and Lassen [25] deal with nondeterminism, and establish soundness in all relevant cases, although full abstraction fails. The first author, together with Alberti, Crubillé and Sangiorgi [10], [8] have studied the probabilistic case, where full abstraction can indeed be obtained if call-by-value evaluation is employed.

X. CONCLUSION

This is the first abstract account on applicative bisimilarity for calculi with effects. The main result is an abstract soundness theorem for a notion of applicative similarity which can be naturally defined as soon as a monad and an associated relator are given which on the one hand serve to give an operational semantics to the algebraic operations, and on the other need to satisfy some mild conditions in order for similarity to be a precongruence. Soundness of bisimilarity is then obtained as a corollary. Many concrete examples are shown to fit into the introduced axiomatics. A notable example is the output monad, for which a definition of applicative similarity based on labeled transition systems as in e.g. [7] is unsound, a fact that the authors discovered after noticing the anomaly, and not vice versa. Nevertheless, it is possible to define a different notion of applicative similarity that fits into our framework and whose associated notion of bisimilarity (Definition 22) coincide with the usual notion of bisimilarity. The reader can consult [9] for details.

A question that we have not addressed in this work, but which is quite natural, is whether an abstract full-abstraction result could exist, analogously to what, e.g., Johann, Simpson, and Voigtländer obtained for their notion of logical relation.

This is a very interesting topic for future work. It is however impossible to get such a theorem without imposing some further, severe, constraints on the class of effects (i.e. monads and relators) of interest, e.g., applicative bisimilarity is well-known not to be fully-abstract in calculi with nondeterministic effects, which perfectly fit in the picture we have drawn in this paper. A promising route towards this challenge would be to understand which class of *tests* (if any) characterise applicative bisimilarity, depending on the underlying monad and relator, this way generalising results by van Breugel, Mislove, Ouaknine and Worrell [40] or Ong [8].

Finally, environmental bisimilarity is known [23] to overcome the limits of applicative bisimilarity in presence of information hiding. Studying the applicability of the methodology developed in this work to environmental bisimilarity is yet another interesting topic for future researches, to which recent work by Sangiorgi and Vignudelli [37] seems to point naturally.

Acknowledgment

The authors would like to thank Raphaëlle Crubillé and the anonymous reviewers for the many useful comments, some of which led to a substantial improvement of our work. Special thanks go to Davide Sangiorgi, Ryo Tanaka, and Valeria Vignudelli for many insightful discussions about the topics of this work. The first and second authors are partially supported by the ANR projects 12ISO2001 PACE, 14CE250005 ELICA, and 16CE250011 REPAS, and by the INRIA-JSPS joint project CRECOGI.

REFERENCES

- [1] Samson Abramsky. The lazy lambda calculus. In D. Turner, editor, *Research Topics in Functional Programming*, pages 65–117. Addison Wesley, 1990.
- [2] Andrew W. Appel and David A. McAllester. An indexed model of recursive types for foundational proof-carrying code. *ACM Trans. Program. Lang. Syst.*, 23(5):657–683, 2001.
- [3] Hendrik P. Barendregt. *The lambda calculus: its syntax and semantics*. Studies in logic and the foundations of mathematics. North-Holland, 1984.
- [4] M. Barr. Relational algebras. *Lect. Notes Math.*, 137:39–55, 1970.
- [5] Nick Benton, Andrew Kennedy, Lennart Beringer, and Martin Hofmann. Relational semantics for effect-based program transformations: higher-order store. In *Proc. of PPDP 2009*, pages 301–312, 2009.
- [6] Ales Bizjak and Lars Birkedal. Step-indexed logical relations for probability. In *Proc. of FOSSACS 2015*, pages 279–294, 2015.
- [7] Roy L. Crole and Andrew D. Gordon. Relating operational and denotational semantics for input/output effects. *Mathematical Structures in Computer Science*, 9(2):125–158, 1999.
- [8] Raphaëlle Crubillé and Ugo Dal Lago. On probabilistic applicative bisimulation and call-by-value λ -calculi. In *Proc. of ESOP 2014*, volume 8410 of *LNCS*, pages 209–228. Springer, 2014.
- [9] Ugo Dal Lago, Francesco Gavazzo, and Paul Levy. Effectful applicative bisimilarity: Monads, relators, and Howe’s method (long version). Available at <http://arxiv.org/abs/1704.04647>, 2017.
- [10] Ugo Dal Lago, Davide Sangiorgi, and Michele Alberti. On coinductive equivalences for higher-order probabilistic functional programs. In *Proc. of POPL 2014*, pages 297–308, 2014.
- [11] Ugo Dal Lago and Margherita Zorzi. Probabilistic operational semantics for the lambda calculus. *RAIRO - Theor. Inf. and Applic.*, 46(3):413–450, 2012.
- [12] Vincent Danos and Russell Harmer. Probabilistic game semantics. *ACM Transactions on Computational Logic*, 3(3):359–382, 2002.
- [13] Brian A. Davey and Hilary A. Priestley. *Introduction to lattices and order*. Cambridge University Press, 1990.
- [14] Ugo de’Liguoro and Adolfo Piperno. Non deterministic extensions of untyped lambda-calculus. *Inf. Comput.*, 122(2):149–177, 1995.
- [15] Joseph A. Goguen, James W. Thatcher, Eric G. Wagner, and Jesse B. Wright. Initial algebra semantics and continuous algebras. *J. ACM*, 24(1):68–95, 1977.
- [16] Sergey Goncharov and Lutz Schröder. A relatively complete generic hoare logic for order-enriched effects. In *Proc. of LICS 2013*, pages 273–282. IEEE Computer Society, 2013.
- [17] Andrew D. Gordon. A tutorial on co-induction and functional programming. In *Workshops in Computing*, pages 78–95. Springer London, September 1994.
- [18] Jean Goubault-Larrecq, Slawomir Lasota, and David Nowak. Logical relations for monadic types. *Mathematical Structures in Computer Science*, 18(6):1169–1217, 2008.
- [19] Douglas J. Howe. Proving congruence of bisimulation in functional programming languages. *Inf. Comput.*, 124(2):103–112, 1996.
- [20] Patricia Johann, Alex Simpson, and Janis Voigtländer. A generic operational metatheory for algebraic effects. In *Proc. of LICS 2010*, pages 209–218. IEEE Computer Society, 2010.
- [21] Claire Jones. *Probabilistic Non-determinism*. PhD thesis, University of Edinburgh, 1990. Available as Technical Report CST-63–90.
- [22] Shin-ya Katsumata and Tetsuya Sato. *Preorders on Monads and Coalgebraic Simulations*, pages 145–160. Springer, 2013.
- [23] Vasileios Koutavas, Paul Blain Levy, and Eijiro Sumii. From applicative to environmental bisimulation. *Electr. Notes Theor. Comput. Sci.*, 276:215–235, 2011.
- [24] Søren B. Lassen. Relational reasoning about contexts. In Andrew D. Gordon and Andrew M. Pitts, editors, *Higher Order Operational Techniques in Semantics*, pages 91–136. 1998.
- [25] Søren B. Lassen. *Relational Reasoning about Functions and Nondeterminism*. PhD thesis, Dept. of Computer Science, University of Aarhus, May 1998.
- [26] Paul Blain Levy. Infinitary Howe’s method. *Electr. Notes Theor. Comput. Sci.*, 164(1):85–104, 2006.
- [27] Paul Blain Levy. Similarity quotients as final coalgebras. In *Proc. of FOSSACS 2011*, volume 6604 of *LNCS*, pages 27–41, 2011.
- [28] Paul Blain Levy, John Power, and Hayo Thielecke. Modelling environments in call-by-value programming languages. *Inf. Comput.*, 185(2):182–210, 2003.
- [29] Saunders MacLane. *Categories for the Working Mathematician*. Springer-Verlag, 1971.
- [30] Eugenio Moggi. Computational lambda-calculus and monads. In *Proc. of (LICS 1989)*, pages 14–23. IEEE Computer Society, 1989.
- [31] J. Morris. *Lambda Calculus Models of Programming Languages*. PhD thesis, MIT, 1969.
- [32] C.-H. Luke Ong. Non-determinism in a functional setting. In *Proc. of LICS 1993*, pages 275–286. IEEE Computer Society, 1993.
- [33] Andrew M. Pitts. Howe’s method for higher-order languages. In D. Sangiorgi and J. Rutten, editors, *Advanced Topics in Bisimulation and Coinduction*, volume 52 of *Cambridge Tracts in Theoretical Computer Science*, chapter 5, pages 197–232. Cambridge University Press, November 2011.
- [34] Gordon D. Plotkin and John Power. Adequacy for algebraic effects. In *Proc. of FOSSACS 2001*, pages 1–24, 2001.
- [35] Gordon D. Plotkin and John Power. Notions of computation determine monads. In *Proc. of FOSSACS 2002*, pages 342–356, 2002.
- [36] Jan J. M. M. Rutten. Universal coalgebra: a theory of systems. *Theor. Comput. Sci.*, 249(1):3–80, 2000.
- [37] Davide Sangiorgi and Valeria Vignudelli. Environmental bisimulations for probabilistic higher-order languages. In *Proc. of POPL 2016*, pages 595–607, 2016.
- [38] Alexander Schrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, Inc., New York, NY, USA, 1986.
- [39] Albert Marchienus Thijs et al. *Simulation and fixpoint semantics*. Rijksuniversiteit Groningen, 1996.
- [40] Franck van Breugel, Michael W. Mislove, Joël Ouaknine, and James Worrell. Domain theory, testing and simulation for labelled markov processes. *Theor. Comput. Sci.*, 333(1-2):171–197, 2005.