

Compositional Expected-Cost Analysis of Functional Probabilistic Programs*

PEDRO H. AZEVEDO DE AMORIM, University of Oxford, UK

Reasoning about the cost of executing programs is one of the fundamental questions in computer science. In the context of programming with probabilities, however, the notion of cost stops being deterministic, since it depends on the probabilistic samples made throughout the execution of the program. This interaction is further complicated by the non-trivial interaction between cost, recursion and evaluation strategy.

In this work we introduce **cert**: a Call-By-Push-Value (CBPV) metalanguage for reasoning about probabilistic cost. We equip **cert** with an operational cost semantics and define two denotational semantics – a cost semantics and an expected-cost semantics. We prove operational soundness and adequacy for the denotational cost semantics and a cost adequacy theorem for the expected-cost semantics.

We formally relate both denotational semantics by stating and proving a novel *effect simulation* property for CBPV. We also prove a canonicity property of the expected-cost semantics as the minimal semantics for expected cost and probability by building on recent advances on monadic probabilistic semantics.

Finally, we illustrate the expressivity of **cert** and the expected-cost semantics by presenting case-studies ranging from randomized algorithms to stochastic processes and show how our semantics capture their intended expected cost.

1 INTRODUCTION

Reasoning about the runtime efficiency of programs is one of the core concepts in computer science. Concepts such as recurrence relations are powerful and flexible tools that are used in modeling and reasoning about the cost structure of programs. One of the advantages of functional programming is that the recursive structure of programs lends itself very well to extracting their cost recurrence relations. However, due to the subtle interaction of cost, higher-order functions and evaluation strategies, one must be careful when trying to establish a formal connection between programs and recurrence relations.

These issues have been explored by Danner et al. [9–11], where the authors, in a span of several papers, have developed semantic techniques that lay on firm ground the cost analysis of functional programs with inductive types. This line of work has culminated in [24], where the authors use a Call-By-Push-Value (CBPV) metalanguage and the writer monad to define a recurrence extraction mechanism for a functional language with recursion and list datatypes.

Though their results are impressive, the only effect their technique can handle is recursion. This is limiting because many problems can be solved more efficiently by having access to probabilistic primitives, so-called randomized algorithms [31] – a notable example being a probabilistic variant of the quicksort algorithm, where by choosing the pivot element uniformly at random, it is possible to average out some of its pathological behavior, such as requiring $O(n^2)$ comparisons when the input list is in reverse order.

When it comes to the mathematical machinery used to reason about these algorithms, familiar tools from deterministic cost analysis such as recurrence relations, have also been successfully developed in the probabilistic setting [21]. Unfortunately, much like in the deterministic case before the work of Danner et al., there has not been a systematic semantic study on how to bring these cost analysis concepts to expressive functional languages.

*For the purpose of Open Access the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

Our Work: Compositional Methods for Probabilistic Cost. In this work we shed some light on the semantic foundations of expected cost analysis in the context of recursive probabilistic functional programs. We start by defining **cert**: a CBPV metalanguage with operations for sampling from uniform distributions and for incrementing the cost of programs. This metalanguage is expressive enough to represent the cost structure of recursive probabilistic algorithms and stochastic processes, and is the first cost-aware metalanguage that can accommodate both continuous distributions and different evaluation strategies.

Besides equipping our metalanguage with an equational theory and an operational cost semantics, we provide two denotational semantics for reasoning about the cost of probabilistic programs. The first one, which we call the *cost semantics*, uses the familiar writer monad transformer to combine a cost monad with a subprobability monad. The second semantics, which we call the *expected cost semantics*, encapsulates the compositional structure of the expected cost as a monad, allowing us to give a denotational semantics that directly tracks the expected cost of programs.

Then, in order to justify the mutual validity of these distinct semantics, we show that the expected cost semantics is a sound approximation to the cost semantics and to the equational theory. In the absence of recursion, we show that this approximation is an equality while it is an upper bound in the presence of unbounded recursion. In order to achieve this soundness result, we state and prove a generalization of the effect simulation problem [23] beyond base types. This generalization interacts well with the CBPV type structure and provides a novel semantic technique for doing relational reasoning of effectful programs.

Importantly, we compare our semantics to the influential pre-expectation program transformations introduced by Kaminski et al. [19]. We argue that our expected cost semantics provides the right abstractions for expected-cost reasoning, since many useful properties that are proved by delicate syntactic arguments in the pre-expected semantics, hold automatically and unconditionally in our semantics. We also prove an unexpected connection between these two monads by showing that the expected cost monad can be obtained as the minimal submonad of the pre-expectation monad that can accommodate probability and expected cost, showing that besides providing a compositional account to expected cost, the expected cost monad is canonical.

As applications, we showcase the capabilities of our semantics by using it to reason about the expected cost of the probabilistic algorithms quicksort and quickselect and stochastic processes such as symmetric random walk. As a guiding example, throughout the paper we will use geometric distributions to illustrate different aspects of **cert** and its semantics.

Our approach contrasts with other work done on expected cost analysis where the language analyzed was either imperative [3, 19, 25] or first-order [26, 36], or the cost structure was given by non-compositional methods [2, 6, 38]. In spirit, the closest to what we have done is [24], which does denotational cost analysis in a deterministic recursive setting, and [1], which uses a continuation-passing style transformation to reason about the expected cost of a call-by-value language.

Our contributions. The main contributions of this paper are the following

- The metalanguage **cert**, a CBPV variant with primitives for increasing cost (charge c) and for sampling from uniform distributions (uniform) (§2)
- A novel expected cost semantics based on an expected cost monad that accommodates familiar and useful reasoning principles (§3.2)
- A generalization of the effect simulation problem for cost semantics using a novel logical relations for denotational relational reasoning (§4.1)
- An operational cost semantics for which we prove the expected cost semantics adequate (§4.3)

$$\begin{aligned}
\bar{\tau} &:= F\tau \mid \tau \rightarrow \bar{\tau} \\
\tau &:= U\bar{\tau} \mid 1 \mid \mathbb{N} \mid \mathbb{R} \mid \tau \times \tau \\
t, u &:= \lambda x. t \mid t V \mid \text{ifZero } V \text{ then } t \text{ else } u \mid \text{force } V \mid (x \leftarrow t); u \\
&\mid \text{produce } V \mid \text{let } x \text{ be } V \text{ in } t \mid \text{succ } t \mid \text{pred } t \\
&\mid \text{let } (x, y) = V \text{ in } t \\
V &:= x \mid () \mid n \in \mathbb{N} \mid r \in \mathbb{R} \mid \text{thunk } t \mid (V_1, V_2) \\
T &:= \text{produce } V \mid \lambda x. t \mid \text{force } x \mid \text{ifZero } x \text{ then } t \text{ else } u
\end{aligned}$$

Fig. 1. Types and Terms of CBPV

$$\begin{array}{c}
\frac{}{\Gamma_1, x : \tau, \Gamma_2 \vdash_v x : \tau} \qquad \frac{n \in \mathbb{N}}{\Gamma \vdash_v n : \mathbb{N}} \qquad \frac{r \in \mathbb{R}}{\Gamma \vdash_v r : \mathbb{R}} \qquad \frac{}{\Gamma \vdash_v () : 1} \\
\\
\frac{\Gamma \vdash_v V : \mathbb{N} \quad \Gamma \vdash_c t : \bar{\tau} \quad \Gamma \vdash_c u : \bar{\tau}}{\Gamma \vdash^c \text{ifZero } V \text{ then } t \text{ else } u : \bar{\tau}} \qquad \frac{\Gamma \vdash_v V_1 : \tau_1 \quad \Gamma \vdash_v V_2 : \tau_2}{\Gamma \vdash_v (V_1, V_2) : \tau_1 \times \tau_2} \\
\\
\frac{\text{op} \in \mathcal{O}(\tau, \tau')}{\Gamma \vdash \text{op} : \tau \rightarrow F\tau'} \qquad \frac{\Gamma, x : \tau \vdash_c t : \bar{\tau}}{\Gamma \vdash_c \lambda x. t : \tau \rightarrow \bar{\tau}} \qquad \frac{\Gamma \vdash_v V : \tau \quad \Gamma \vdash_c t : \tau \rightarrow \bar{\tau}}{\Gamma \vdash_c t V : \bar{\tau}} \\
\\
\frac{\Gamma \vdash_v V : \tau}{\Gamma \vdash_c \text{produce } V : F\tau} \qquad \frac{\Gamma \vdash_c t : \bar{\tau}}{\Gamma \vdash_v \text{thunk } t : U\bar{\tau}} \qquad \frac{\Gamma \vdash_v V : U\bar{\tau}}{\Gamma \vdash_c \text{force } V : \bar{\tau}} \\
\\
\frac{\Gamma \vdash_c t : F\tau' \quad \Gamma, x : \tau' \vdash_c u : \bar{\tau}}{\Gamma \vdash_c (x \leftarrow t); u : \bar{\tau}} \qquad \frac{\Gamma \vdash_v V : \tau_1 \times \tau_2 \quad \Gamma, x : \tau_1, y : \tau_2 \vdash_c t : \bar{\tau}}{\Gamma \vdash_c \text{let } (x, y) = V \text{ in } t : \bar{\tau}}
\end{array}$$

Fig. 2. CBPV typing rules

- A universal property of the expected cost monad as the minimal submonad of the continuation monad that can accommodate subprobability distributions and cost (§5)

We also justify the applicability of our semantics through use-cases that illustrate how the expected cost semantics can be used to reason about expected cost of stochastic processes and randomized algorithms.

2 cert : A PROBABILISTIC COST-AWARE METALANGUAGE

In this section we introduce the type system, equational theory and operational semantics of **cert**. The language is a Call-By-Push-Value (CBPV) [27] calculus extended with operations for cost, probabilistic sampling and recursion. The operational semantics is defined as a weighted Markov chain that accounts for the cost and output distributions of programs. It is defined using the standard measure-theoretic treatment of operational semantics for continuous distributions, cf. Vákár et al. [37].

We chose CBPV as the core of **cert** due to its type and term-level separation of values and computations. Such a separation allows for a fine grained control over the execution of programs,

providing a uniform treatment of different evaluation strategies such as Call-By-Name (CBN) and Call-By-Value (CBV).

Figure 1 depicts the CBPV syntax. Note that the base types 1 , \mathbb{N} and \mathbb{R} are value types, the product types is also a value and arrow types are computation types that receive a value type as input and a computation type as output. At the center of the CBPV formalism are the type constructors F and U which allows types to move between value types and computation types. The constructor F plays a similar role to the monadic type constructor T from the monadic λ -calculus [30], while U is used to represent suspended – or thunked – computations.

This two-level approach also manifests itself at the type judgement level, where the judgement $\Gamma \vdash_v V : \tau$ is only defined for value types, while $\Gamma \vdash_c t : \bar{\tau}$ is defined for computation types, as shown in Figure 2. Both contexts only bind values, which justifies the arrow type having a value type in its domain, so that lambda abstractions only introduce values to the context. The if-then-else operation checks if the guard V is 0, in which case it returns the first branch, and otherwise it returns the second branch. The language is also parametric on a set of operations O which contains arithmetic functions. The product introduction rule pairs two values while its elimination rule unpairs a product and uses them in a computation. Lambda abstraction binds a new value to the context while application applies a function to a value.

The less familiar rules are those for the type constructors F and U . The introduction rule for computations is produce V , which is the computation that does not incur any effect and just outputs the value V , while the introduction rule for U , thunk t , suspends the computation t . Its elimination rule force V resumes the suspended computation V . The last rule, $x \leftarrow t; u$ is what makes it possible to chain effectful computations together, since it receives a computation of type $F\tau$ as input, runs it and binds the result to the continuation u , which eventually will output a computation of type $\bar{\tau}$. This is a generalization of the monadic let rule where the output type does not have to be of type $F\tau$.

The syntax differs a bit from the monadic semantics of effects, but every strong monad over a Cartesian closed category can interpret the CBPV calculus, as we describe in Appendix A.

Though this language is effective as a core calculus, by itself it cannot do much, since it has no “native” effect operations, meaning that there are no programs with non-trivial side-effects. In this section we extend CBPV so that it can program with three different effects: cost, probability and unbounded recursion. We call this extension **cert**, for **calculus for expected run time**, and we conclude the section by presenting its equational theory and operational semantics.

2.1 Cost and Probabilistic Effects

As it is common in denotational approaches to cost semantics, it is assumed that there is a cost monoid \mathbb{C} – usually interpreted by \mathbb{N} and addition – which acts on programs by operations charge c that increases the current cost of the computation by c units, for every $c : \mathbb{C}$. The value types are extended with a type \mathbb{C} and constants $\cdot \vdash_v c : \mathbb{C}$. Furthermore, since we also want to program with probabilities and unbounded recursion, we extend the language with a sampling primitive, as well as recursive definitions:

$$\frac{\Gamma \vdash_v V : \mathbb{C}}{\Gamma \vdash_c \text{charge } V : F1} \qquad \frac{}{\Gamma \vdash_c \text{uniform} : F\mathbb{R}} \qquad \frac{\Gamma, x : U\bar{\tau} \vdash_c t : \bar{\tau}}{\Gamma \vdash_c \text{fix } x. t : \bar{\tau}}$$

The operation `uniform` uniformly samples a real number from the interval $[0, 1]$ and `fix` is the familiar fixed-point operator used for defining recursive programs. In interest of reducing visual pollution and simplifying the presentation, `charge V ; t` desugars to `($x \leftarrow \text{charge } V$); t` , when x is not used in the body of t , and we will assume that the cost monoid is \mathbb{N} .

<pre> fix f : list(τ) \rightarrow FN. λl : list(τ). case l of nil \Rightarrow produce 0 (hd, tl) \Rightarrow $n \leftarrow$ (force f) tl produce (1 + n) </pre>	<pre> fix f : list(τ) \rightarrow F(list(τ) \times list(τ)). λl : list(τ) . λp : $\tau \rightarrow$ FN. case l of nil \Rightarrow produce (nil, nil) (hd, tl) \Rightarrow $n \leftarrow p$ hd (l_1, l_2) \leftarrow (force f) p tl if n then produce (cons hd l_1, l_2) else produce ($l_1, \text{cons hd } l_2$) </pre>
---	--

Fig. 3. Length function (left) and filter function (right).

With the uniform distribution primitive it is possible to define uniform distributions over discrete sets which, given a natural number n , outputs a uniform distribution $\text{rand } n$ over the set $\{0, \dots, n-1\}$. This can be desugared to the program $\lambda n. x \leftarrow \text{uniform}; \text{produce } (\lfloor nx \rfloor) : \mathbb{N} \rightarrow \text{FN}$, where $\lfloor \cdot \rfloor$ is the floor function. Biased coins \oplus_p desugar to $\lambda p. x \leftarrow \text{uniform}; x \leq p : \mathbb{R} \rightarrow \text{FN}$, where $\leq : \mathbb{R} \rightarrow \mathbb{R} \rightarrow \text{FN}$ is the comparison function that returns 0 if the first argument is less or equal to the second argument and 1 otherwise.

Example 2.1 (Geometric distribution). With these primitives we can already program non-trivial distributions. For instance, the geometric distribution can be expressed as the program

$$\cdot \vdash_c \text{fix } x. (\text{produce } 0) \oplus_{0.5} ((y \leftarrow \text{force } x); \text{produce } (1 + y)) : \text{FN},$$

Operationally, the program flips a fair coin, if the output is 0, it outputs 0, otherwise it recurses on x , binds the value to y and outputs $1 + y$. By the typing rule of recursive definitions, the variable x is a thunk, meaning that it must be forced before executing it.

By having fine-grained control over which operations have a cost, it is possible to orchestrate your program with charge c operations in order to encode different cost models. For instance, if we want to keep track of how many coins were tossed when running the geometric distribution, we can modify it as such

$$\text{fix } x. \text{charge } 1; (\text{produce } 0) \oplus_{0.5} (y \leftarrow \text{force } x; \text{produce } (1 + y)) : \text{FN}$$

Example 2.2 (Deterministic Programs). The charge operation can also be used to keep track of the number of recursive calls in your program. For instance, a recursive program that computes the factorial function can be instrumented to count the number of recursive calls as follows:

$$\text{fix } f. \lambda n. \text{ifZero } n \text{ then } (\text{produce } 0) \text{ else } (\text{charge } 1; n * (\text{force } f)(n - 1))$$

Whenever the if-guard is false, the cost is incremented by 1 and the function is recursively called.

2.2 Lists

Frequently, cost analysis are defined for algorithms defined over inductive data types, such as lists. As such, we will also extend our language with lists over value types.

$$\begin{aligned}
\tau &:= \dots \mid \text{list}(\tau) \\
V &:= \dots \mid \text{nil} \mid \text{cons } V_1 V_2 \\
t &:= \dots \mid (\text{case } x \text{ of nil} \Rightarrow t \mid \text{cons } x \text{ xs} \Rightarrow u) \\
T &:= \dots \mid (\text{case } x \text{ of nil} \Rightarrow t \mid \text{cons } x \text{ xs} \Rightarrow u)
\end{aligned}$$

$$\begin{array}{c}
\frac{}{\Gamma \vdash^v \text{nil} : \text{list}(\tau)} \qquad \frac{\Gamma \vdash^v V_1 : \tau \quad \Gamma \vdash^v V_2 : \text{list}(\tau)}{\Gamma \vdash^v \text{cons } V_1 V_2 : \text{list}(\tau)} \\
\\
\frac{\Gamma \vdash^v V : \text{list}(\tau) \quad \Gamma \vdash^c t : \bar{\tau} \quad \Gamma, x : \tau, xs : \text{list}(\tau) \vdash^c u : \bar{\tau}}{\Gamma \vdash^c \text{case } V \text{ of nil} \Rightarrow t \mid \text{cons } x \text{ xs} \Rightarrow u : \bar{\tau}}
\end{array}$$

The primitive nil is the empty list, cons appends a value to the front of a list and case is for pattern-matching on lists and, in the presence of fix, can be used for defining non-structurally recursive functions over lists.

Example 2.3. The functions that computes the length of a list and a binary version of the familiar filter function that outputs two lists, one for the true elements and one for the false elements, are, respectively, defined in the left and right parts of Figure 3. Note that since we have adopted a \mathbb{N} -valued if-statement, the predicate p above outputs a natural number. Furthermore, since the recursion operation adds a thunk to the context, in order to call the recursive function you must first force its execution.

Example 2.4. In Figure 4 we have defined a randomized version of the quicksort algorithm that counts the number of comparisons done. In in, we are accessing the r -th element of a list l using the familiar syntax $l[r]$. The algorithm is very similar to the non-randomized quicksort with the exception of choosing the pivot element with the command $r \leftarrow \text{rand } \text{len}$ that uniformly chooses an element from the set $\{0, \dots, |l|\}$, where $|-\|$ is the length function.

We conclude this section by mentioning that there are many other sensible extensions, such as recursive and sum types. For our purposes, they are not necessary and so, in order to keep the language simple, we omit them. That being said, from a semantic point of view, these extensions are well-understood and straightforward to be accommodated by the denotational semantics we present in Section 3.

2.3 Equational Theory

We want to define a syntactic sound approximation to the expected cost of programs. We do this by extending the usual equational theory of CBPV with rules for the monoid structure of the charge operation. We present some of the equational theory in Figure 5, with other rules which are standard in CBPV languages shown in Appendix A. The first two rules are the familiar β and η -rules for the arrow type, the 0MON and ACTMON rules are the monoid equations for the charge operation. The rule THUNKFORCE says that forcing a thunked computation is the same thing as running the computation, the rules IFZ and IFS explain how if-statements interact with natural

```

fix f : list( $\mathbb{N}$ )  $\rightarrow$  F(list( $\mathbb{N}$ )).
 $\lambda l$  : list( $\mathbb{N}$ ).
case l of
| nil  $\Rightarrow$ 
  produce nil
| (hd, tl)  $\Rightarrow$ 
  len  $\leftarrow$  length l
  r  $\leftarrow$  rand len
  pivot  $\leftarrow$  l[r]
  ( $l_1, l_2$ )  $\leftarrow$  biFilter ( $\lambda n$ . charge 1;  $n \leq$  pivot) l
  less  $\leftarrow$  force f l1
  greater  $\leftarrow$  force f l2
  produce (less  $\#$  pivot :: greater)

```

Fig. 4. Randomized quicksort

numbers and the rule `Fix` is the fixed point equation that unfolds one recursive call of the recursive computation t .

2.4 Operational Semantics

Since we are interested in modeling the cost of running programs, we will define an operational cost semantics which is closer to the execution model of programs. When defining semantics for probabilistic languages with continuous distributions, one must be careful to define it so that it is a measurable function.

In this section, we begin by showing the the `cert` syntax can be equipped with a measurable space structure that makes the natural syntax operations, such as substitution, measurable. Then, the operational semantics can be defined as a Markov kernel over the syntax, as it is usually done in probabilistic operational semantics for continuous distributions [12, 37]. After defining the operational semantics, we conclude by proving the subject reduction property.

Syntax Spaces and Kernels. Before defining the operational semantics, we need some definitions from measure theory.

Definition 2.5. A measurable space is a pair $(X, \Sigma_X \subseteq \mathcal{P}(X))$, where X is a set and Σ is a σ -algebra, i.e. a collection of subsets that contains the empty set and is closed under complements and countable union.

Definition 2.6. A measurable function $f : (X, \Sigma_X) \rightarrow (Y, \Sigma_Y)$ is a function $f : X \rightarrow Y$ such that for every $A \in \Sigma_Y$, $f^{-1}(A) \in \Sigma_X$.

Definition 2.7. A subprobability distribution over a measurable space (X, Σ_X) is a function $\mu : \Sigma_X \rightarrow [0, 1]$ such that $\mu(\emptyset) = 0$, $\mu(X) \leq 1$ and $\mu(\bigcup_{n \in \mathbb{N}} A_n) = \sum_{n \in \mathbb{N}} \mu(A_n)$.

The operational semantics will be modeled as a (sub)Markov kernel, a generalization of transition matrices and Markov chains.

$$\begin{array}{c}
\beta\text{-LAW} \\
\frac{\Gamma, x : \tau \vdash_c t : \bar{\tau} \quad \Gamma \vdash_\sigma V : \tau}{\Gamma \vdash t\{V/x\} = (\lambda x. t) V : \bar{\tau}}
\end{array}
\quad
\begin{array}{c}
\eta\text{-LAW} \\
\frac{\Gamma \vdash_c t : \tau \rightarrow \bar{\tau}}{\Gamma \vdash_c (\lambda x. t x) = t : \tau \rightarrow \bar{\tau}}
\end{array}
\quad
\begin{array}{c}
0\text{MON} \\
\frac{\Gamma \vdash_c t : \bar{\tau}}{\Gamma \vdash (\text{charge } 0; t) = t : \bar{\tau}}
\end{array}$$

$$\begin{array}{c}
\text{ACTMON} \\
\frac{}{\Gamma \vdash \text{charge } c; \text{charge } d = \text{charge } c + d : F1}
\end{array}
\quad
\begin{array}{c}
\text{THUNKFORCE} \\
\frac{\Gamma \vdash_c t : \bar{\tau}}{\Gamma \vdash \text{force } (\text{thunk } (t)) = t : \bar{\tau}}
\end{array}$$

$$\begin{array}{c}
\text{IFZ} \\
\frac{\Gamma \vdash_c t : \bar{\tau} \quad \Gamma \vdash_c u : \bar{\tau}}{\Gamma \vdash \text{ifZero } 0 \text{ then } t \text{ else } u = t : \bar{\tau}}
\end{array}
\quad
\begin{array}{c}
\text{IFS} \\
\frac{\Gamma \vdash_c t : \bar{\tau} \quad \Gamma \vdash_c u : \bar{\tau}}{\Gamma \vdash \text{ifZero } (n + 1) \text{ then } t \text{ else } u = u : \bar{\tau}}
\end{array}$$

$$\begin{array}{c}
\text{FIX} \\
\frac{\Gamma, x : U\bar{\tau} \vdash_c t : \bar{\tau}}{\Gamma \vdash (\text{fix } x. t) = t\{x/\text{thunk } (\text{fix } x. t)\} : \bar{\tau}}
\end{array}$$

Fig. 5. Equational Theory (Selected Rules)

Definition 2.8. A subMarkov kernel between measurable spaces (X, Σ_X) and (Y, Σ_Y) is a function $f : X \times \Sigma_Y \rightarrow [0, 1]$ such that

- For every $x : X$, $f(x, -) : \Sigma_Y \rightarrow [0, 1]$ is a subprobability distribution
- For every $A : \Sigma_Y$, $f(-, A) : X \rightarrow [0, 1]$ is a measurable function

We denote the set of computation terms by Λ , the set of values by \mathcal{Val} and the set of terminal computations by T . Let t (resp. V) be a computation (resp. value), fix the term traversal order left-to-right and let $z_1, z_2, \dots, z_n, \dots$ be a sequence of distinct and ordered variables disjoint from the set of term variables. The traversal order gives rise to a canonical enumeration of t 's (resp. V 's) occurrences of numerals $r \in \mathbb{R}$, which we denote by the sequence r_1, r_2, \dots, r_n . By substituting these occurrences by the variables z_1, z_2, \dots, z_n , we obtain the term $t\{z_1, \dots, z_n/r_1, \dots, r_n\}$ (resp. $V\{z_1, \dots, z_n/r_1, \dots, r_n\}$). Let Λ_n (resp. \mathcal{Val}_n) be the set of such substituted terms with exactly n numerals.

Note that the sets Λ_n and \mathcal{Val}_n are countable and that there are bijections $\Lambda \cong \sum_{n:\mathbb{N}, t:\Lambda_n} \mathbb{R}^n$ and $\mathcal{Val} \cong \sum_{n:\mathbb{N}, t:\mathcal{Val}_n} \mathbb{R}^n$, where Σ is the dependent sum operation: for instance, every computation term t can be decomposed into a sequence of its numerals r_1, \dots, r_n and substituted term $t\{z_1, \dots, z_n/r_1, \dots, r_n\}$, and, conversely, every sequence of numerals and substituted term t can be mapped to the term $t\{r_1, \dots, r_n/z_1, \dots, z_n\}$ – note the reversed order of substitution.

Therefore, we can equip Λ with the coproduct σ -algebra: $(\Lambda, \Sigma_\Lambda) = \sum_{n:\mathbb{N}, t:\Lambda_n} (\mathbb{R}^n, \Sigma_{\mathbb{R}^n})$. More concretely, a subset $A \subset \Lambda$ is measurable if, and only if,

$$\forall n \in \mathbb{N}, t : \Lambda_n, \{(r_1, \dots, r_n) \in \mathbb{R}^n \mid t\{r_1, \dots, r_n/z_1, \dots, z_n\} \in A\} \in \Sigma_{\mathbb{R}^n}$$

The measurable space structure of \mathcal{Val} is defined using a similar coproduct.

Given a pair of a context Γ and a computation type $\bar{\tau}$, the measurable spaces $\Lambda^{\Gamma \vdash \bar{\tau}}$ and $T^{\Gamma \vdash \bar{\tau}}$ are the subspaces of well-typed computations and terminal computations under context Γ and output type $\bar{\tau}$, respectively. Given a value type τ , the measurable space $\mathcal{Val}^{\Gamma \vdash \tau} \subseteq \mathcal{Val}$ is the subspace of well-typed values under context Γ and output type τ . The measurable space structures of $\Lambda^{\Gamma \vdash \bar{\tau}}$, T , $T^{\Gamma \vdash \bar{\tau}}$ and $\mathcal{Val}^{\Gamma \vdash \tau}$ are defined using the appropriate subspace σ -algebras. The following metatheoretic lemmas are useful and standard.

Lemma 2.9. *If $\Gamma, x : \tau \vdash_c t : \bar{\tau}$ and $\Gamma \vdash_v V : \tau$ then $\Gamma \vdash_c t\{V/x\} : \bar{\tau}$.*

Lemma 2.10. *For every variable x , the substitution function $\cdot\{ \cdot / x \} : \Lambda \times \mathcal{Val} \rightarrow \Lambda$ is measurable.*

Therefore, for every context Γ , computation type $\bar{\tau}$ and value type τ , the substitution function restricts to measurable functions $\cdot\{ \cdot / x \} : \Lambda^{\Gamma, x : \tau \vdash \bar{\tau}} \times \mathcal{Val}^{\Gamma \vdash \tau} \rightarrow \Lambda^{\Gamma \vdash \bar{\tau}}$.

Operational Kernels. In order to capture the cost of running a computation, we define *costful* kernels as a subMarkov kernel $X \times \Sigma_{\mathbb{N} \times Y} \rightarrow [0, 1]$. Given two costful kernels f and g , their composition $f \circ g : X \times \Sigma_{\mathbb{N} \times Z} \rightarrow [0, 1]$ is defined, with slight abuse of notation, as:

$$(g \circ f)(x, n_1 + n_2, C) = \int_Y g(y, n_2, C) f(x, n_1, -) (dy)$$

In plain terms, every term reduces to a subprobability distribution over costs and terminal computations. Their composition is defined so that the probability that the composition cost is $n_1 + n_2$ is the probability that the input f will cost n_1 and the continuation g will cost n_2 , i.e. the product of both events averaged out using an integral. We use the Haskell syntax $\gg=$ for kernel composition.

We can now define the operational semantics as the limit of a sequence of approximate semantics given by costful kernels $\Downarrow_n : \Lambda \times \Sigma_{\mathbb{N} \times T} \rightarrow [0, 1]$. The approximate semantics \Downarrow_n are defined by recursion on n , where the base case is defined as $\Downarrow_0 = \perp$, i.e. the 0 measure. When $n > 0$, the recursive definition is depicted in Figure 6.

Though we are using the familiar relational definition of operational semantics, they are functional in nature. We use the notation of Vákár et al. [37], where the inference rule

$$\frac{k_1(t)w_1 \quad k_2(t, w_1)w_2 \quad \dots \quad k_n(t, w_1, \dots, w_n)v}{l(t)f(t, w_1, \dots, w_n, v)}$$

denotes the kernel $k_1(t) \gg= (\lambda w_1. k_2(t, w_1) \gg= \dots \delta_{f(t, w_1, \dots, w_n, v)})$. We also simplify the presentation by using *guarded* kernel composition. For instance, in the β -reduction rule, whenever t reduces to something which is not a λ -abstraction, the kernel composition loops. Besides the non-standard presentation, the semantics is a fairly standard CBPV big-step semantics [27]. For example, terminal computations cannot reduce any further, so they output a pointmass distribution over 0 cost and themselves. The non-standard rules are the effectful operations, where the charge operation steps to $()$ while increasing the cost of the computation; the sample operation reduces to an independent distribution of the pointmass distribution at 0 and the Lebesgue uniform measure λ on the interval $[0, 1]$.

It follows by a simple induction that the semantics \Downarrow_n is monotonic in n . Since the space of subprobability distributions forms a CPO, we can define the semantics as the supremum of its finite approximations $\Downarrow = \bigsqcup_n \Downarrow_n$. As usual, it is possible to prove subject reduction by induction on well-typed terms.

Lemma 2.11 (Subject reduction). *If $\Gamma \vdash_c t : \bar{\tau}$, then the composition $\Downarrow \circ \iota : \Lambda^{\Gamma \vdash \bar{\tau}} \rightarrow P_{\leq 1}(\mathbb{N} \times T)$ factors as $\Lambda^{\Gamma \vdash \bar{\tau}} \rightarrow P_{\leq 1}(\mathbb{N} \times T^{\Gamma \vdash \bar{\tau}}) \hookrightarrow P_{\leq 1}(\mathbb{N} \times T)$, where $\iota : \Lambda^{\Gamma \vdash \bar{\tau}} \rightarrow \Lambda$ is the inclusion function. More colloquially, well-typedness is stable under the operational semantics.*

3 DENOTATIONAL SEMANTICS

This section presents two concrete denotational semantics to our language:

- A cost semantics that serves as a denotational baseline for compositionally computing the cost distribution of probabilistic programs.

$$\begin{array}{c}
\frac{}{\text{produce } V \Downarrow_n \delta_{(0, \text{produce } V)}} \quad \frac{t \Downarrow_n \mu}{\text{force } (\text{thunk } t) \Downarrow_n \mu} \quad \frac{}{\text{uniform } \Downarrow_n \delta_0 \otimes \lambda} \quad \frac{}{\lambda x. t \Downarrow_n \delta_{(0, \lambda x. t)}} \\
\frac{t \Downarrow_n \lambda x. u \quad u\{V/x\} \Downarrow_{n-1} \mu}{t V \Downarrow_n \mu} \quad \frac{}{\text{charge } r \Downarrow_n \delta_{(r, \text{produce } ())}} \\
\frac{t \Downarrow_n \text{produce } V \quad u\{V/x\} \Downarrow_{n-1} \mu}{(x \leftarrow t); u \Downarrow_n \mu} \quad \frac{t\{\text{thunk fix } x. t/x\} \Downarrow_{n-1} \mu}{\text{fix } x. t \Downarrow_n \mu} \quad \frac{t \Downarrow_n \mu}{\text{ifZero } 0 \text{ then } t \text{ else } u \Downarrow_n \mu} \\
\frac{u \Downarrow_n \mu}{\text{ifZero } (n+1) \text{ then } t \text{ else } u \Downarrow_n \mu} \quad \frac{t\{V_1, V_2/x_1, x_2\} \Downarrow_n \mu}{\text{let } (x_1, x_2) = (V_1, V_2) \text{ in } t \Downarrow_{n-1} \mu} \\
\frac{t \Downarrow_{n-1} \mu}{\text{case nil of nil } \Rightarrow t \mid \text{cons } x \text{ xs } \Rightarrow u \Downarrow_n \mu} \quad \frac{t\{V_1, V_2/x, \text{xs}\} \Downarrow_{n-1} \mu}{\text{case } (\text{cons } V_1 V_2) \text{ of nil } \Rightarrow t \mid \text{cons } x \text{ xs } \Rightarrow u \Downarrow_n \mu}
\end{array}$$

Fig. 6. Big-Step Operational Semantics

- An expected cost semantics that, while it cannot reason about as many quantitative properties of cost as the cost semantics, such as tail-bounds and higher-moments, it provides a compositional account to the *expected cost*.

Both semantics will be defined over the category $\omega\mathbf{Qbs}$ of ω -quasi Borel spaces [37], a Cartesian closed category that admits a probabilistic powerdomain of subprobability distributions $P_{\leq 1}$. By using the writer monad transformer $P_{\leq 1}(\mathbb{C} \times -)$, it can also accommodate cost operations, as we explain in Section 3.1. With this monad it is possible to define the expected cost to be the expected value of the cost distribution $P_{\leq 1}(\mathbb{C})$.

Unfortunately, this approach is non-compositional. In order to compute the expected cost of a program of type $F\tau$, we must first compute its (compositional) semantics which can then be used to obtain a distribution over the cost and then apply the expectation formula to it. We work around this issue by constructing a novel expected cost monad in Section 3.3 that makes the expected cost a part of the semantics and, as such, it is compositionally computed. We start this section by going over important definitions and constructions for $\omega\mathbf{Qbs}$, we then define the cost semantics, followed by the expected cost semantics.

3.1 ω -quasi Borel spaces

We now introduce the semantic machinery used in the interpretation of **cert**. Due to requirement of higher-order functions, probability and unbounded recursion, we are somewhat limited in terms of which semantic domain to use. We use the category of ω -quasi Borel spaces, a domain-theoretic version of quasi Borel spaces [15].

Definition 3.1 ([37]). An ω -quasi Borel space is a triple (X, \leq, M_X) such that, (X, \leq) is a ω -complete partial order (ωCPO), i.e. it is a partial order closed under suprema of ascending sequences, and $M_X \subseteq \mathbb{R} \rightarrow X$ is the set of *random elements* with the following properties:

- All constant functions are in M_X
- If $f : \mathbb{R} \rightarrow \mathbb{R}$ is a measurable function and $p \in M_X$, then $p \circ f \in M_X$

- If $\mathbb{R} = \bigcup_{n \in \mathbb{N}} U_n$, where for every n , U_n are pairwise-disjoint and Borel-measurable, and $\alpha_n \in M_X$ then the function $\alpha(x) = \alpha_n(x)$ if, and only if, $x \in U_n$ is also an element of M_X .
- For every ascending chain $\{f_n\}_n \subseteq M_X$, i.e. for every $x \in \mathbb{R}$, $f_n(x) \leq f_{n+1}(x)$, the pointwise supremum $\bigsqcup_n f_n$ is in M_X .

Note that, in the definition above, ω CPOs do not assume the existence of a least element, e.g. for every set X , the discrete poset $(X, =)$ is an ω CPO.

Definition 3.2. A measurable function between ω -quasi Borel spaces is a Scott continuous function $f : X \rightarrow Y$ — i.e. preserves suprema of ascending chains — such that for every $p \in M_X$, $f \circ p \in M_Y$.

Definition 3.3. The category $\omega\mathbf{Qbs}$ has ω -quasi Borel spaces as objects and measurable functions as morphisms.

Theorem 3.4 ([37]). *The category $\omega\mathbf{Qbs}$ is Cartesian closed.*

Furthermore, there is a full and faithful functor $\mathbf{Meas} \rightarrow \omega\mathbf{Qbs}$. More concretely, if you interpret a program that has as inputs and output measurable spaces, its denotation in $\omega\mathbf{Qbs}$ will be a measurable function, even if the program uses higher-order functions, and any measurable function could potentially be the denotation of the program.

Inductive types. As shown in previous work [37], $\omega\mathbf{Qbs}$ can also soundly accommodate full recursive types. In particular, it can give semantics to lists over A by solving the domain equation $\text{list}(A) \cong 1 + A \times \text{list}(A)$.

It is convenient that in $\omega\mathbf{Qbs}$, the set of lists over A with appropriate random elements and partial order is a solution to the domain equation and it is the smallest one, i.e. it is an initial algebra. This means that when reasoning about lists expressed in `cert`, you may assume that they are just the set of lists over sets.

Probability and Partiality Monads. It is possible to construct probabilistic powerdomains in $\omega\mathbf{Qbs}$, making it possible to use this category as a semantic basis for languages with probabilistic primitives. Furthermore, the ω CPO structure can also be used to construct a partiality monad, making it possible to give semantics to programs with unbounded recursion. We are assuming familiarity with basic concepts from category theory such as monads and use the notation $(T, \eta^T, (-)_T^\#)$, where T is an endofunctor, $\eta^T : 1 \rightarrow T$ and $(-)_T^\# : (- \Rightarrow T) \rightarrow (T- \Rightarrow T)$ are the unit and bind natural transformations, respectively. The monad is said to be *strong* if there is a natural transformation $st : - \times T \Rightarrow T(- \times -)$ making certain diagrams commute [30]. When it is clear from the context, we will simply write η and $(-)_\#$, without the sub and superscript, respectively.

Lemma 3.5 ([37]). *The category $\omega\mathbf{Qbs}$ admits strong commutative monads P and $P_{\leq 1}$ of probability and sub-probability distributions, respectively.*

Categorically, $P_{\leq 1}$ is defined as a submonad of the continuation monad $(- \rightarrow [0, \infty]) \rightarrow [0, \infty]$ and its monad structure is similar to the one from probability monads in \mathbf{Meas} , i.e. the unit at a point $a : A$ is given by the point mass distribution δ_a and $f^\#(\mu)$ is given by integrating f over the input distribution μ . A more detailed presentation of this construction will be given in Section 5.

Furthermore, by construction, $\omega\mathbf{Qbs}$ admits a morphism $\int_A : (P_{\leq 1}A) \times (A \rightarrow \{0, 1\}) \rightarrow [0, 1]$ that maps a subprobability distribution and a “measurable set” of A into its measure. For example, if A is a measurable space, for every measurable set $X : A \rightarrow \{0, 1\}$, and for every subprobability distribution $\mu : P_{\leq 1}A$, the map $(\mu, X) \mapsto \mu(X)$ is an $\omega\mathbf{Qbs}$ morphism and is equal to \int_A .

As we have mentioned above, it is also possible to define a lifting monad in $\omega\mathbf{Qbs}$ that adds a least element \perp to a space, making them *pointed* ω CPOs. This monad, combined with the ω CPO

$$\llbracket \text{charge } c \rrbracket_{CS} = \delta_{(c,0)} \quad \llbracket \text{uniform} \rrbracket_{CS} = \delta_0 \otimes \lambda \quad \llbracket \text{fix } x. t \rrbracket_{CS} = \bigsqcup_n \llbracket t \rrbracket_{CS}^n (\perp)$$

Fig. 7. Cost semantics of operations

structure, is used to guarantee the existence of fixed points of endomorphisms between pointed ω CPOs.

Definition 3.6 ([37]). The lifting monad in $\omega\mathbf{Qbs } X_\perp$ adds a fresh element \perp to X , makes it the least element and the random elements M_{X_\perp} are the functions $f : \mathbb{R} \rightarrow X_\perp$ such that there is a Borel measurable set \mathcal{B} and a map $\alpha : \mathbb{R} \rightarrow X$ in M_X such that $f(x) = \alpha(x)$ for $x \in \mathcal{B}$ and \perp otherwise.

The machinery we have defined so far is expressive enough to interpret **cert**, with exception of its cost operations. In non-effectful languages, the writer monad $(\mathbb{C} \times -)$ can be used to give semantics to cost operations such as charge c .

Definition 3.7. If $(\mathbb{C}, 0, +)$ is a monoid, then $\mathbb{C} \times -$ is a monad – the *writer* monad – where the unit at a point a is $(0, a)$ and given a morphism $f : A \rightarrow \mathbb{C} \times B$, $f^\#(c, a) = (c + (\pi_1 \circ f)(a), (\pi_2 \circ f)(a))$, where $\pi_i : A_1 \times A_2 \rightarrow A_i$ is the i -th projection.

What follows is how to combine the non-probabilistic cost monad $(\mathbb{C} \times -)$ with $P_{\leq 1}$ in order to define a probabilistic cost semantics.

3.2 A probabilistic cost semantics

Contrary to the deterministic case, the cost of a probabilistic computation is not a single value, it is a distribution over costs. For instance, consider the program:

$$\cdot \vdash_c (\text{charge } 1; \text{produce } 0) \oplus (\text{produce } 2) : FN$$

it either returns 2 without costing anything, or it returns 0 with a cost of 1. Denotationally, this program should be the distribution $\frac{1}{2}(\delta_{(1,0)} + \delta_{(0,2)})$. With equal probability, the program will either cost 1 and output 0 or cost 0 and output 2.

In the deterministic case, it is possible to encode the cost at the semantic-level by using the *writer* monad $\mathbb{C} \times -$. For probabilistic cost-analysis we can use the writer monad transformer.

Lemma 3.8. *If $T : C \rightarrow C$ is a strong monad then $T(\mathbb{C} \times -)$ is a strong monad.*

PROOF. The strength of a monad is a natural transformation $A \times TB \rightarrow T(A \times B)$. When instantiating A to be \mathbb{C} , we can conclude that there is a distributive law between the writer monad and T , which allows us to conclude that $T(\mathbb{C} \times -)$ is a monad. Its strength is defined as $st^T; T(st^{\mathbb{C} \times -}) : A \times T(\mathbb{C} \times B) \rightarrow T(A \times (\mathbb{C} \times B)) \rightarrow T(\mathbb{C} \times (A \times B))$. \square

When instantiating T to be the subprobability monad $P_{\leq 1}$, we get a monad for probabilistic cost, which justifies the denotation of the program $(\text{charge } 1; \text{produce } 0) \oplus (\text{produce } 2)$ being a distribution of a pair of a cost and natural number.

By using the monadic semantics of CBPV, we get a cost-aware probabilistic semantics, where most of its definitions follow the standard monadic CBPV semantics shown in Appendix A – denoted as $\llbracket \cdot \rrbracket_{CS}^v$ for values and $\llbracket \cdot \rrbracket_{CS}^c$ for computations. The noteworthy interpretations are for the effectful operations, whose semantics are depicted in Figure 7, and for the cost monoid, which is interpreted as the additive natural numbers $(\mathbb{N}, 0, +)$.

With this semantics, we now define the expected cost of a distribution:

Definition 3.9. Let $\mu : P_{\leq 1}[0, \infty]$, its expected value is $\mathbb{E}(\mu) = \int_{[0, \infty]} x \, d\mu$.

In the definition above we have chosen the most general domain for \mathbb{E} , but for every measurable subset $X \subseteq [0, \infty]$ the expected distribution formula can be restricted to distributions over X . In particular, it is possible to restrict this function to have $P_{\leq 1}(\mathbb{N})$ as its domain.

Example 3.10 (Geometric Distribution). This semantics makes it possible to reason about the geometric distribution defined as the program $^1 \cdot \vdash \text{fix } x.0 \oplus (1 + x) : F\mathbb{N}$. It is possible to show that this program indeed denotes the geometric distribution by unfolding the semantics and obtaining the fixed point equation $\mu = \frac{1}{2}(\delta_0 + P_{\leq 1}(\lambda x.1 + x)(\mu))$, for which the geometric distribution is a solution.

Since we are interested in reasoning about the cost of programs, it is possible to reason about the expected amount of coins flipped during its execution by adding the charge 1 operation as follows $\text{fix } x. \text{charge}_1; (0 \oplus (1 + x)) : F\mathbb{N}$, i.e. whenever a new coin is flipped, as modeled by the \oplus operation, the cost increases by one. By construction, the cost distribution will also follow a geometric distribution.

If we want to compute the actual expected value, we must compute $\sum_{n:\mathbb{N}} \frac{n}{2^n}$. This particular infinite sum can be calculated by using a standard trick. In the next section we show how to encode this trick in the semantics itself, significantly simplifying the computation of the expected value.

Expected cost and non-termination. When designing metalanguages for reasoning about cost of programs, it is expected that the cost of a non-terminating computation is infinite. Therefore, since subprobability distributions of mass less than 1 model possibly non-terminating computation, their cost should be ∞ . This can be achieved by modifying the expected cost function to $\infty \cdot (1 - \mu(\mathbb{C})) + \int_{\mathbb{C}} x \, d\mu$ and defining $\infty \cdot 0 = 0$, meaning that every computation with non-zero chance of termination would, by definition, have infinite expected cost.

There are a few problems with this solution. Since the soundness proof of Section 4 requires the expected cost to be a morphism $P_{\leq 1}(\mathbb{C}) \rightarrow [0, \infty]$ in $\omega\mathbf{Qbs}$, we have to equip the extended positive real line with a set of random elements and a partial order. Traditionally, the bottom element of $[0, \infty]$ would be interpreted as ∞ , in line with the slogan that “no information” equals infinite cost – this is the approach taken by [24]. However, as we will see in Section 6, when we have recursively defined programs of type $F\tau$, their expected costs are given by the supremum of increasing sequences of real numbers, starting from 0. Suggesting that 0 needs to be the bottom element while ∞ is the top element.

This alternative is also problematic, unfortunately, since it is not Scott-continuous. As an example consider the geometric distribution, whose cost distribution is given by the supremum of the sequence $\{\sum_{i=1}^n 2^{-i} \delta_i\}_n$. Therefore, $\mathbb{E}(\sum_{i=1}^n 2^{-i} \delta_i) = \infty$, for every $n \in \mathbb{N}$, resulting in

$$\sup_n \mathbb{E} \left(\sum_{i=1}^n 2^{-i} \delta_i \right) = \infty \neq \mathbb{E} \left(\sup_n \sum_{i=1}^n 2^{-i} \delta_i \right) = \mathbb{E}(\text{geom}) = 2,$$

where *geom* is the geometric distribution. Other approaches have not dealt with these difficulties because, unlike other effects, probabilities present genuinely interesting recursive programs of type $F\mathbb{N}$, as illustrated by the geometric distribution itself. In contrast, in [24], by a monotonicity argument, every recursive program of type $F\mathbb{N}$ either diverges or outputs a constant. Therefore, assigning infinite cost to these programs, though not very realistic from a modeling point of view, is not too damaging to the semantics.

¹For the sake of simplicity we have elided the some of the bureaucracy of CBPV, such as produce and force .

$$\llbracket \text{charge } c \rrbracket_{EC} = (c, \delta_{()}) \quad \llbracket \text{uniform} \rrbracket_{EC} = (0, \lambda) \quad \llbracket \text{fix } x. t \rrbracket = \bigsqcup_n \llbracket t \rrbracket_{EC}^n (\perp)$$

Fig. 8. Expected cost semantics of operations

With this in mind, we argue that Definition 3.9 is a sensible choice. In Section 4, we go over the consequences of this definition insofar as the cost semantics interacts with the expected cost semantics.

3.3 A semantics for expected cost

The semantics just proposed can compositionally compute the cost distribution of programs, but compositionality is broken when computing its expected cost. Indeed, after computing the distribution, we must compute the expected cost of an arbitrarily complex distribution. We fix this by defining a novel expected cost monad.

We achieve this by defining a monad structure on the functor $[0, \infty] \times P_{\leq 1}$: every computation will be denoted by an extended positive real number, i.e. its expected cost, and a subprobability distribution over its output. The monad's unit at a point $a : A$ is the pair $(0, \delta_a)$ and the bind operation $(-)^{\#}$ adds the expected cost of the input with the average of the expected cost of the output. Formally, given an $\omega\mathbf{Qbs}$ morphism $f : X \rightarrow [0, \infty] \times P_{\leq 1}Y$, its bind is the function $f^{\#}(r, \mu) = (r + \int (\pi_1 \circ f) d\mu, (\pi_2 \circ f)_{P_{\leq 1}}^{\#}(\mu))$.

Theorem 3.11. *The triple $([0, \infty] \times P_{\leq 1}, \eta, (-)^{\#})$ is a strong monad.*

PROOF. The proof can be found in Appendix F. □

With this monad it is possible to define a new semantics to **cert** that interprets the effectful operations a bit differently from the cost semantics, as we depict in Figure 8, where $\llbracket \cdot \rrbracket_{EC}^c$ is the computation semantics while $\llbracket \cdot \rrbracket_{CS}^v$ is the value semantics; the cost monoid is still interpreted as \mathbb{N} .

Example 3.12 (Revisiting the geometric distribution). Unfolding the semantics $\llbracket \text{geom} \rrbracket_{EC}$ gives us the fixed point equation $E = 1 + (1 - \frac{1}{2})E$, i.e. $E = 2$. This can be readily generalized to arbitrary $p \in [0, 1]$, giving the equation $E_p = 1 + (1 - p)E_p$.

As we have noted in the previous section, the cost semantics can be used to reason about the expected cost by using Definition 3.9. Something which will play an important role in our soundness proof is the fact that this definition interacts well with the monadic structure of $P_{\leq 1}$.

Lemma 3.13. *Let $\mu : P_{\leq 1}A$ and $f : A \rightarrow P_{\leq 1}([0, \infty])$, $\mathbb{E}(f^{\#}(\mu)) = \int_A \mathbb{E}(f(a))\mu(da)$.*

PROOF. This can be proved by unfolding the definitions

$$E(f^{\#}(\mu)) = \int_{[0, \infty]} x \left(\int_A f(a)\mu(da) \right) (dx) = \int_A \int_{[0, \infty]} xf(a)(dx)\mu(da) = \int_A \mathbb{E}(f(a))\mu(da)$$

Note that in the third equation we had to reorder the integrals, which is valid because $P_{\leq 1}$ is commutative. □

With this lemma in mind, we state some basic definitions and lemmas that allows us to describe precisely how the cost and expected cost semantics relate.

Definition 3.14. A monad morphism is a natural transformation $\gamma : T \rightarrow S$, where $(T, \eta^T, (-)^{\#}_T)$ and $(S, \eta^S, (-)^{\#}_S)$ are monads over the same category, such that $\gamma \circ \eta^T = \eta^S$ and $(\gamma \circ g)^{\#}_S \circ \gamma = \gamma \circ g^{\#}_T$, for every $g : A \rightarrow TB$.

Theorem 3.15. *There is a monad morphism $E : P(\mathbb{N} \times -) \rightarrow [0, \infty] \times P$.*

PROOF. We define the morphism $E_A(\mu) = (\mathbb{E}(P(\pi_1)(\mu)), P(\pi_2)(\mu))$. The first monad morphism equation follows by inspection and the second one follows mainly from Lemma 3.13, when restricting it to the probabilistic distributions, i.e. total mass equal to 1. \square

Lemma 3.16. *The natural transformation E , when extend to subprobability distributions, is not a monad morphism.*

PROOF. Let $\frac{1}{2}(\delta_{(0,1)} + \delta_{(1,2)})$ be a distribution over $\mathbb{C} \times \mathbb{N}$ and $f(0) = \frac{1}{2}\delta_0$, $f(n+1) = 0$ be a subprobability kernel. It follows by inspection that $((E \circ f)^\#_{[0, \infty] \times P_{\leq 1}} \circ E)(\mu) \neq (E \circ f^\#_{P_{\leq 1}(\mathbb{N} \times -)})(\mu)$ \square

Theorem 3.15 says that the different cost semantics interact well in the probabilistic case. In the subprobabilistic case this is not true, as illustrated by Lemma 3.16. This formalizes the intuitions behind the subtleties in the interaction of expected cost and non-termination explained in the previous section.

4 SOUNDNESS THEOREMS

We have proposed four ways of reasoning about expected cost: by using the equational theory, the operational semantics or by using either of the denotational semantics. Something which is not clear at first is the extent of how much these semantics are reasoning about the same property. We begin this section by proving soundness properties of the expected cost semantics with respect to the denotational cost semantics by stating and proving a generalization of the *effect simulation problem*. Due to subtle interactions between cost, probability and non-termination, we provide separate analyses for the probabilistic and subprobabilistic cases.

In order to justify that the denotational semantics is reasoning about an operational notion of cost, we prove standard soundness and adequacy results of the denotational cost semantics with respect to the operational cost semantics. We also show that both cost and expected cost semantics are sound with respect to the equational theory.

4.1 Denotational Soundness

Denotational Soundness: Probabilistic Case. We start our investigation by restricting the semantic monads to their submonads of probability distributions, i.e. we use the monads $P(\mathbb{C} \times -)$ and $[0, \infty] \times P$, and remove the recursion operation, while keeping the semantics otherwise unchanged.

The soundness property we are interested in is the following: given a closed program $\cdot \vdash_c t : F\tau$, then the expected value for the second marginal of $\llbracket t \rrbracket_{CS}$ is equal to $\pi_1(\llbracket t \rrbracket_{EC})$. In the literature, similar properties have been called the *effect simulation problem* and many semantic techniques have been developed for solving it [23]. Unfortunately, even one of the most general ones, called $\top\top$ -lifting [22], does not meet our requirements.

A simplified version of the main theorem in [23] is the following:

Theorem 4.1 (Th. 7 of [23]). *Let C be a category, $T : C \rightarrow C$ and $S : C \rightarrow C$ be monads such that there is a monad morphism $\gamma : T \rightarrow S$, then for every program $\cdot \vdash_c t : F\mathbb{N}$, $\gamma_{\mathbb{N}}(\llbracket t \rrbracket_{CS}^c) = \llbracket t \rrbracket_{EC}^c$.*

We can see how this theorem is relevant to what we want to prove: there are two monads $P(\mathbb{C} \times -)$ and $[0, \infty] \times P$, a monad morphism E between them and we want to show that their output distributions are the same and the expected cost of both semantics are equal.

The problem with the theorem above is that, even in its most general form proved by Katsumata [23], it can only handle base types and it gives no guarantees for programs of type, say, $\text{list}(\mathbb{N}) \rightarrow F(\text{list}(\mathbb{N}))$, which is a significant limitation to the case studies we study in Section 6.

We circumvent these issues by defining a two-level logical relations inspired by the $\top\top$ -lifting construction. The two-level structure mimics the value/computation types distinction present in CBPV and gives stronger reasoning principles than in the work of Katsumata [23]. At the core of our argument is the relational-lifting construction, defined as follows:

Definition 4.2. Let $\mathcal{R} \subseteq A \times B$ be a complete binary relation, i.e. it is closed under suprema of ascending sequences, its lifting $\mathcal{R}^\# \subseteq P_{\leq 1}(A) \times P_{\leq 1}(B)$ is defined as $\mu \mathcal{R}^\# \nu$ if there is a distribution $\theta : P_{\leq 1}(\mathcal{R})$ such that its first and second marginals are, respectively, μ and ν .

This definition interacts well with the monadic structure of $P_{\leq 1}$. Concretely, it is stable with respect to the unit and bind of $P_{\leq 1}$. This definition can be restricted to the P monad. We now define our logical relations as two families of relations, one for value types and one for computation types:

$$\begin{aligned} \mathcal{V}_\tau &\subseteq \llbracket \tau \rrbracket_{CS}^v \times \llbracket \tau \rrbracket_{EC}^v & \mathcal{C}_{\bar{\tau}} &\subseteq \llbracket \bar{\tau} \rrbracket_{CS}^c \times \llbracket \bar{\tau} \rrbracket_{EC}^c \\ \mathcal{V}_{\mathbb{N}} &= \{(n, n) \mid n \in \mathbb{N}\} & \mathcal{C}_{F\tau} &= \{((r, \nu), \mu) \mid \mathbb{E}(\mu_1) = r \wedge \nu \mathcal{V}_\tau^\# \mu_2\} \\ \mathcal{V}_{U\bar{\tau}} &= \mathcal{C}_{\bar{\tau}} & \mathcal{C}_{\tau \rightarrow \bar{\tau}} &= \{(f_1, f_2) \mid \forall x_1, x_2, x_1 \mathcal{V}_\tau x_2 \Rightarrow f_1(x_1) \mathcal{C}_{\bar{\tau}} f_2(x_2)\} \\ \mathcal{V}_{\tau_1 \times \tau_2} &= \mathcal{V}_{\tau_1} \times \mathcal{V}_{\tau_2} \\ \mathcal{V}_{\text{list}(\tau)} &= \text{list}(\mathcal{V}_\tau) \end{aligned}$$

In order to prove the fundamental theorem of logical relations, we first show the following lemmas, where the first one follows by induction:

Lemma 4.3. For every τ (resp. $\bar{\tau}$), the relation \mathcal{V}_τ (resp. $\mathcal{C}_{\bar{\tau}}$) is an ω CPO, where the partial order structure is the same as the one from $\llbracket \tau \rrbracket_{CS}^v \times \llbracket \tau \rrbracket_{EC}^v$ (resp. $\llbracket \bar{\tau} \rrbracket_{CS}^c \times \llbracket \bar{\tau} \rrbracket_{EC}^c$). Furthermore, the computation relations have a least element.

Lemma 4.4. For every type τ (resp. $\bar{\tau}$), there is a set M and partial order \leq such that the triple $(\mathcal{V}_\tau, M, \leq)$ (resp. $(\mathcal{C}_{\bar{\tau}}, M, \leq)$) is an ω -quasi Borel space such that the injection function is a morphism in ω Qbs.

PROOF. We only make explicit the proof for value types, since the case of computation types is basically the same. We define the order \leq to be the same as the one in $\llbracket \tau \rrbracket_{CS}^v \times \llbracket \tau \rrbracket_{EC}^v$ and M to be the restricted random elements $\{f \in M_{\tau_1} \mid f(\mathbb{R}) \subseteq \mathcal{V}_\tau\}$.

Since by the lemma above the logical relations are ω CPOs, M is closed under suprema of ascending chains, and $(\mathcal{V}_\tau, M, \leq)$ is an ω -quasi Borel space. The injection into $\llbracket \tau \rrbracket_{CS}^v \times \llbracket \tau \rrbracket_{EC}^v$ being a morphism follows by construction. \square

We now state the denotational soundness theorem:

Theorem 4.5. For every $\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n, \Gamma \vdash_v V : \tau, \Gamma \vdash_c t : \bar{\tau}$ and if for every $1 \leq i \leq n, \cdot \vdash_v V_i : \tau_i$ and $\llbracket V_i \rrbracket_{CS} \mathcal{V}_{\tau_i} \llbracket V_i \rrbracket_{EC}$, then

$$\begin{aligned} &\left[\text{let } \bar{x}_i = \bar{V}_i \text{ in } t \right]_{CS}^c \mathcal{C}_{\bar{\tau}}^c \left[\text{let } \bar{x}_i = \bar{V}_i \text{ in } t \right]_{EC}^c \text{ and} \\ &\left[\text{let } \bar{x}_i = \bar{V}_i \text{ in } V \right]_{CS}^v \mathcal{V}_\tau \left[\text{let } \bar{x}_i = \bar{V}_i \text{ in } V \right]_{EC}^v, \end{aligned}$$

where the notation $\bar{x}_i = \bar{V}_i$ means a list of n let-bindings or, in the case of values, a list of substitutions.

PROOF. The proof follows by induction and is shown in Appendix B. \square

We now have a very precise sense in which the expected-cost semantics is related to the cost semantics:

Corollary 4.6. *The expected-cost semantics is sound with respect to the cost semantics, i.e. for every program $\cdot \vdash_c t : F\tau$, the expected cost of the second marginal of $\llbracket t \rrbracket_{CS}^c$ is equal to $\pi_1(\llbracket t \rrbracket_{EC}^c)$.*

That being said, the theorem above is only valid for probability distributions. In the subprobabilistic case only a weaker version of this theorem holds.

Denotational Soundness: Subprobabilistic Case. Programming without recursion is somewhat limiting. Without an infinitely supported base distribution, every definable distribution has finite support. In particular, it would not be possible to define the geometric distribution.

Therefore, it seems reasonable to prove the soundness theorem above for $P_{\leq 1}$. Unfortunately, it does not hold in the subprobabilistic case, as alluded to in Lemma 3.16.

Example 4.7. Consider the programs

$$\begin{aligned} t &= \text{charge } 2; \text{produce } 0 \\ u &= \lambda x. \text{ifZero } x \text{ then } (\perp \oplus (\text{charge } 4; \text{produce } 0)) \text{ else } \perp \end{aligned}$$

We can show $E(\llbracket x \leftarrow t; u \ x \rrbracket_{CS}^c) = (3, \frac{1}{2}\delta_0) \neq (4, \frac{1}{2}\delta_0) = \llbracket x \leftarrow t; u \ x \rrbracket_{EC}^c$

Even though the counterexample above invalidates our soundness theorem, it is still possible to prove a weaker variant of Corollary 4.6. This is achieved by using essentially the same logical relations as before, with the exception of $C_{F\tau}$, which now becomes

$$C_{F\tau} = \{((r, v), \mu) \mid \mathbb{E}(\mu_1) \leq r \wedge v \mathcal{V}_\tau^\# \mu_2\}$$

Corollary 4.8. *The expected-cost semantics is sound with respect to the cost semantics, i.e. for every program $\cdot \vdash_c t : F\tau$, the expected cost of the second marginal of $\llbracket t \rrbracket_{CS}^c$ is at most $\pi_1(\llbracket t \rrbracket_{EC}^c)$.*

PROOF. The proof can also be found in Appendix B. □

4.2 Equational Soundness

In the last section we have argued that in the presence of unbounded recursion, the cost semantics has some undesirable properties which are not present in the expected cost semantics. That being said, when it comes to the base equational theory of Figure 5, both semantics are sound with respect to it, showing that there is still some harmony between them.

Theorem 4.9. *If $\Gamma \vdash_c t = u : \bar{\tau}$ then $\llbracket t \rrbracket_{EC}^c = \llbracket u \rrbracket_{EC}^c$ and $\llbracket t \rrbracket_{CS}^c = \llbracket u \rrbracket_{CS}^c$.*

PROOF. The proof follows by induction on the equality rules, where the inductive cases follow directly from the inductive hypothesis while the base cases follow by inspection. For instance, the equation $\text{charge } 0; t = t$ is true because \mathbb{N} is a monoid and 0 is its unit. Once again, the full equational theory is shown in Figure 11. □

It is useful to understand the extent in which these equational theories differ. For instance, the cost semantics validates the equation $\perp; t = \perp = t; \perp$. This equation is too extreme for the purposes of expected cost, since it says that $\text{charge } c; \perp = \perp$. An even more egregious equation that it satisfies is $\text{fix } x. \text{charge } 1; x = \perp$. That equation says that the cost of infinity is the same as no cost at all, as long as the program does not terminate.

These equations are connected to the commutativity equation:

$$\frac{\Gamma \vdash_c x : F\tau_1 \quad \Gamma \vdash_c u : F\tau_2 \quad \Gamma, x : \tau_1, y : \tau_2 \vdash_c t' : \bar{\tau}}{\Gamma \vdash_c (x \leftarrow t; y \leftarrow u; t') = (y \leftarrow u; x \leftarrow t; t') : \bar{\tau}}$$

Theorem 4.10 (c.f. Appendix F). *The cost semantics validates the commutativity equation.*

This equation is usually useful for reasoning about probabilistic programs. However, it is too strong for the purposes of reasoning about expected cost. Indeed, consider the programs

$$\begin{aligned} t &= \perp; \text{charge } c; \text{produce } () \\ u &= \text{charge } c; \perp; \text{produce } () \end{aligned}$$

From an operational point of view, the first program will run a non-terminating program and never reach the charge c operation, while the second one starts by increasing the cost by c and then loops forever.

When it comes to modeling the cost of real programs, these two programs should not be the same since, for instance, if the charge operation is modeling a monetary cost, such as a call to an API, only the second program will cost something. Fortunately, the expected cost monad is not commutative, making it a more physically-justified monad.

Lemma 4.11. *The monad $[0, \infty] \times P_{\leq 1}$ is not commutative.*

PROOF. When $c > 0$, the terms above are a counterexample: $\llbracket t \rrbracket_{EC} = (0, 0) \neq (c, 0) = \llbracket u \rrbracket_{EC}$ \square

4.3 Operational Soundness and Adequacy

We conclude this section by proving some metatheoretic properties of the operational semantics and show how it relates to the denotational cost semantics. The main results of this section are the soundness and adequacy properties of the operational semantics with respect to the cost semantics and the cost adequacy theorem of the operational semantics with respect to the expected cost semantics.

A consequence of Lemma 2.11 is that it becomes possible to compose the operational and denotational semantics and obtain morphisms $\llbracket \Downarrow \rrbracket^{\Gamma+\bar{\tau}} : \Lambda^{\Gamma+\bar{\tau}} \rightarrow P_{\leq 1}(\mathbb{N} \times (\llbracket \Gamma \rrbracket_{CS} \rightarrow \llbracket \bar{\tau} \rrbracket_{CS}))$. In particular, for closed programs, $\llbracket \Downarrow \rrbracket^{+\bar{\tau}} : \Lambda^{+\bar{\tau}} \rightarrow P_{\leq 1}(\mathbb{N} \times (\llbracket \bar{\tau} \rrbracket_{CS}))$. We now state the soundness theorem.

Theorem 4.12 (Soundness). *For every closed computation $\cdot \vdash_c t : \bar{\tau}$, $\alpha_{\bar{\tau}}(\llbracket \Downarrow (t) \rrbracket_{CS}) \leq \llbracket t \rrbracket_{CS}$.*

PROOF. As usual, since the operational semantics is defined as the supremum of \Downarrow_n , the proof follows by induction on n and t . The proof can be found in Appendix C. \square

The next step is proving the adequacy theorem, which in this context is equivalent to the converse of the soundness theorem above. We prove it by first defining a pair of logical relations $\triangleright_{\bar{\tau}}$ and $\blacktriangleleft_{\bar{\tau}}$.

$$\begin{aligned} \triangleright_{\bar{\tau}} &\subseteq \mathcal{V}al^{+\circ\bar{\tau}} \times \llbracket \bar{\tau} \rrbracket_{CS} & \blacktriangleleft_{\bar{\tau}} &\subseteq \Lambda^{+\circ\bar{\tau}} \times \llbracket \bar{\tau} \rrbracket_{CS} \\ n \triangleright_{\mathbb{N}} n &\iff \top & t \blacktriangleleft_{F\bar{\tau}} \mu &\iff \mu \leq \llbracket \Downarrow t \rrbracket \\ r \triangleright_{\mathbb{R}} r &\iff \top & t \blacktriangleleft_{\tau \rightarrow \bar{\tau}} f &\iff \forall V a, V \triangleright_{\tau} a \implies (t V) \blacktriangleleft_{\bar{\tau}} f(a) \\ (V_1, V_2) \triangleright_{\tau_1 \times \tau_2} (x, y) &\iff (V_1 \triangleright_{\tau_1} x) \wedge (V_2 \triangleright_{\tau_2} y) \\ l \triangleright_{\text{list}(\bar{\tau})} x &\iff l \text{list}(\triangleright_{\bar{\tau}}) x \\ (\text{thunk } t) \triangleright_{U\bar{\tau}} x &\iff t \blacktriangleleft_{\bar{\tau}} x \end{aligned}$$

Theorem 4.13 (Adequacy). *For every closed computation $\cdot \vdash_c t : \bar{\tau}$, $\llbracket t \rrbracket_{CS} \leq \alpha_{\bar{\tau}}(\llbracket \Downarrow (t) \rrbracket_{CS})$.*

PROOF. The proof follows by a logical relations argument and can be found in Appendix D \square

Corollary 4.14 (Cost Adequacy). *For every closed computation $\cdot \vdash_c t : F\tau$, $\pi_1(\llbracket t \rrbracket_{EC}) \leq \mathbb{E}(P_{\leq 1}(\pi_1)(\llbracket \Downarrow (t) \rrbracket_{CS}))$.*

PROOF. This theorem is a direct consequence of the adequacy theorem above and Corollary 4.8. \square

5 PRE-EXPECTATION SOUNDNESS

We conclude the technical development of the semantics by proving a canonicity property of the expected cost monad. We achieve this by relating our semantics to the work on pre-expectation semantics for expected cost analysis advocated for in previous work [1, 19]. This line of work show that it is possible to reason about the expected cost of programs by using a continuation-passing style transformation and the monad $(X \rightarrow [0, \infty]) \rightarrow [0, \infty]$, hereafter referred to as $K_{[0, \infty]}$.

One of the major strengths of using continuation monads for modeling effects is that they are extremely flexible and, therefore, can encode a wide variety of impure computations by choosing appropriate response types. Unfortunately, when using such general semantics, one loses the structures and invariants of "tailor made" semantics, making reasoning about programs less direct.

In fact, indirections caused by continuation semantics are well-known in the compiler literature, where the administrative reductions added by continuation-passing style makes it harder to recover information from the source program. Indeed, some authors [8, 29] advocate against the use of continuation-passing style precisely because of this obfuscation.

In the context of expected-cost analysis, one such indirection can be seen in Avanzini et al. [3], where the soundness of their cost analysis hinges on proving that the interpretation of programs can be rewritten as the sum of their expected cost and an integration operator — i.e. a (sub)probability distribution. Their proof relies on syntactic invariants of their language and does not hold for the entire continuation semantics, making it brittle when it comes to language extensions. In contrast, our semantics validates a similar decomposition property by definition of the expected cost monad.

In this section we show that the expected cost semantics “lives inside” the pre-expectation semantics without collapsing any equalities. More concretely, we show that there is an injection from the expected cost monad into the continuation monad and that, in a precise sense, the expected cost monad is the minimal submonad of the continuation monad that can accommodate cost and subprobability distributions. This demonstrates the canonicity of our semantics. We achieve this by using the McDermott and Kammar theory of factorization of monad morphisms [20].

We begin by reviewing the theory of factorization systems for monad morphisms and go over how such a structure is used to construct the statistical powerdomain in ω Qbs. We then show how this construction can be extended to the expected cost monad. As an application, we conclude this section by providing a novel and more robust proof of Lemma 7.1 in [3].

Monad Structures and Factorization Systems. Factorization systems capture classes of morphisms that can uniquely factor any morphism. The most familiar example is that of injections and surjections, where every function can be uniquely factored as the composition of a surjection followed by an injection. However, in order to properly generalize this idea, one needs to be a bit more careful:

Definition 5.1 ([5]). A (orthogonal) factorization system over a category C is a pair $(\mathcal{E}, \mathcal{M})$, where both \mathcal{E} and \mathcal{M} are classes of morphisms of C such that

- Every isomorphism belongs to both \mathcal{E} and \mathcal{M}
- Both classes are closed under composition

- For every $e : \mathcal{E}$, $m : \mathcal{M}$, and morphisms f and g there is a unique h making the following diagram commute:

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ e \downarrow & \dashv\!\! \dashv \! \dashrightarrow h & \downarrow m \\ C & \xrightarrow{g} & D \end{array}$$

- Every morphism f in \mathcal{C} can be factored as $f = m \circ e$, where $m : \mathcal{M}$ and $e : \mathcal{E}$.

Note that under these axioms, the factorization is guaranteed to be unique up-to isomorphism.

Example 5.2. As mentioned above, in the category **Set**, the pair $(\{f \mid f \text{ is surjective}\}, \{f \mid f \text{ is injective}\})$ is a factorization system. Given a function $f : A \rightarrow B$, its factorization is $A \rightarrow f(A) \rightarrow B$, where $f(A)$ is the image of f and the function $f(A) \rightarrow B$ is the set inclusion.

Example 5.3 ([20]). In the category **CPO**, the pair $(\{f \mid f \text{ has dense image}\}, \{f \mid f \text{ is order-reflecting}\})$ is a factorization system.

A natural extension of the factorization system above can be defined for the category $\omega\mathbf{Qbs}$. We now recall some definitions and a theorem from [20] that play important role in how the probability monad in $\omega\mathbf{Qbs}$ is defined.

Definition 5.4. A monad structure is an endofunctor S equipped with the monad natural transformations η and $(-)^{\#}$ but without the necessity of satisfying the monad laws.

Definition 5.5. A monad structure morphism between monad structures S_1 and S_2 over the same categories is a natural transformation $S_1 \rightarrow S_2$ such that the monad morphism laws hold.

Theorem 5.6 (Th. 2.5 of [20]). *Let $(\mathcal{E}, \mathcal{M})$ be a factorization structure, S a monad structure, T a monad and $\gamma : S \rightarrow T$ a premonad morphism. If the factorization system is closed under S then γ can be factored as $S \rightarrow M \rightarrow T$, where M is a monad, both morphisms are premonad morphisms and each component of these morphisms are elements of \mathcal{E} and \mathcal{M} .*

This theorem was used by Vákár et al. [37] when defining their statistical powerdomain as the factor of a monad morphism into the pre-expectation monad $K_{[0, \infty]}$. Intuitively, the monad structures we are interested in are those where the monad laws hold up-to an equivalence relation and the factored monads are, roughly, the quotiented monad structures and satisfy a minimality universal property given by the uniqueness of factorizations.

Canonicity. We conclude this section by showing that the expected cost monad is the minimal monad that can accommodate cost and subprobability distributions. We extend the analysis done by Vákár et al. [37] on using Theorem 5.6 to define and prove the minimality of their statistical powerdomain. We now review its construction and then apply it to our semantics.

The “random elements” functor $\mathbb{R} \rightarrow -_{\perp}$ can be equipped with a monad structure similar to, but distinct from, the reader monad, as explained in Section 4.2 of [37]. The measures monad in $\omega\mathbf{Qbs}$ is then defined using the following lemma:

Lemma 5.7 (Sec. 4.2 [37]). *There is a premonad morphism $(\mathbb{R} \rightarrow -_{\perp}) \rightarrow K_{[0, \infty]}$.*

The measure powerdomain T is defined as the factorization of the morphism above. The component of the natural transformation above maps a pair $(f : \mathbb{R} \rightarrow A_{\perp}, g : A \rightarrow [0, \infty])$ to $\int_{\text{dom}(f)} (g \circ f) d\lambda$, where λ is the Lebesgue uniform measure and $\text{dom}(f)$ is the domain of f . The probability powerdomain can be defined as those measures that have total mass 1 or, more formally, as the equalizer $PX \rightarrow TX \rightrightarrows [0, \infty]$, between the constant function $\mu \mapsto 1$ and the total mass

function $\mu \mapsto \int_X d\mu$. The subprobability monad is then defined as $P_{\leq 1}X = P(X_{\perp})$, where the mass of \perp corresponds to the probability of non-termination.

In analogy with the previous results, we can now prove a similar lemma for the expected cost measure monad.

Lemma 5.8. *There is a monad structure on the functor $[0, \infty] \times (\mathbb{R} \rightarrow -_{\perp})$ and a monad structure morphism $\gamma : [0, \infty] \times (\mathbb{R} \rightarrow -_{\perp}) \rightarrow K_{[0, \infty]}$.*

PROOF. There are similarities between the expected cost monad and the monad structure which is given by $\eta(a) = (0, \lambda r. a)$ and $f^{\#}(r, g) = \left(r + \int_{\text{dom}(g)} (\pi_1 \circ f \circ g) d\lambda, (\pi_2 \circ f)_{\mathbb{R} \rightarrow -_{\perp}}^{\#}(g) \right)$. The components of the monad structure morphism are defined as $\gamma_A(r, f) = \lambda k. r + \int (k \circ f) d\lambda$. The proof that this is indeed a monad morphism follows by unfolding the definitions and Lemma 5.7. \square

By Theorem 5.6, and the fact that the factorization system in $\omega\mathbf{Qbs}$ is stable under products, the premonad morphism γ above can be factored as $[0, \infty] \times (\mathbb{R} \rightarrow -_{\perp}) \xrightarrow{\psi} M \rightarrow K_{[0, \infty]}$. By extending the definition of $[0, \infty] \times P$ to unnormalized measures, $[0, \infty] \times T$ is also a monad and we show that it is isomorphic to the factored monad M .

We now show the sense in which $[0, \infty] \times T$ is canonical.

Theorem 5.9. *The factor M is isomorphic to the monad $[0, \infty] \times T$.*

PROOF. The proof can be found in Appendix F. \square

The canonicity of the expected cost monad is given by it being the smallest submonad of the continuation monad that can accommodate expected cost and subprobability distributions. A consequence of this fact is the following:

Lemma 5.10. *There is a monad morphism $\varphi : [0, \infty] \times P_{\leq 1} \rightarrow K_{[0, \infty]}$ which is component-wise order-reflecting.*

PROOF. The monad morphism is the composition $[0, \infty] \times P_{\leq 1} \hookrightarrow [0, \infty] \times T \hookrightarrow K_{[0, \infty]}$ which, by stability under composition, is order-reflecting. \square

Intuitively, while the (sub)probability monad corresponds to the linear continuation monad, the expected cost monad corresponds to the affine continuation monad. An application of the results of this section is that it is possible to recover a purely denotational proof of Lemma 7.1 of [3] that is more robust with respect to language extensions.

Theorem 5.11. *For every program $\cdot \vdash_c t : F\mathbb{R}$, $\llbracket t \rrbracket_{pre} = \varphi(\llbracket t \rrbracket_{EC})$. More explicitly, $\llbracket t \rrbracket_{pre}(f) = \pi_1(\llbracket t \rrbracket_{EC}) + \int f d(\pi_2(\llbracket t \rrbracket_{EC}))$*

PROOF. The proof is a direct application of Theorem 7 in Katsumata [23]. \square

Furthermore, by using the fact that the monad morphism is order-reflecting and, therefore, an injection, we can also prove that if $\llbracket t_1 \rrbracket_{pre} = \llbracket t_2 \rrbracket_{pre}$, then $\llbracket t_1 \rrbracket_{EC} = \llbracket t_2 \rrbracket_{EC}$. It would be interesting to prove a more general version of this theorem using ideas from the denotational soundness proof of Section 4. We leave such generalizations to future work.

6 EXAMPLES

In this section we will show how the expected cost semantics can be used to reason about the expected cost of probabilistic programs. We present four examples, two randomized algorithms and two recursive stochastic processes, illustrating the versatility of *cert*. Due to space constraints, another example can be found in Appendix E.

6.1 Expected coin tosses

A classic problem in basic probability theory is computing the expected number of coin flips necessary in order to obtain n heads in a row. We can model this stochastic process as the following recursive probabilistic program:

```

fix f : ℕ → F1. λn : ℕ.
  if n then
    produce ()
  else
    (force f) (n - 1);
    charge 1;
    (produce () ⊕ (force f) n)

```

For every n , the program above simulates the probabilistic structure of flipping coins until obtaining n heads in a row. When its input is 0, it outputs $()$ without flipping any coins. If the input is greater than 0, in order to flip n heads in a row it must first flip $n - 1$ heads in a row – hence the call to $f(n - 1)$ – flip a new coin while increasing the current counter by 1 and, if it is heads, you have obtained n heads in a row and may output $()$, otherwise you must recursively start the process again from n : the left and right branches of \oplus , respectively.

By unfolding the semantics, we obtain that the expected number of coin tosses is given by the following recurrence relation:

$$T(0) = 0$$

$$T(n + 1) = 1 + T(n) + \frac{1}{2}T(n + 1)$$

Which has the closed-form solution $T(n) = 2(2^n - 1)$.

6.2 Randomized Quicksort

In Example 2.4 we present a program that implements randomized quicksort. If we simply interpret the expected cost of this program denotationally, it will be a function mapping lists to real numbers. This is not how such an analysis is done in practice, where people reason about how the cost increases as the length of the list increases, regardless of which elements it contains.

In our semantics, the denotation of the program is hiding the fact that its cost only depends on the length of its argument. We make this precise by defining a measurable function using the program in Figure 9 $qck_{\mathbb{N}} : \mathbb{N} \rightarrow \mathbb{R} \times P_{\leq 1}\mathbb{N}$ that corresponds to the quicksort structure assuming that the input is a natural number. Let $len : list(\mathbb{N}) \rightarrow \mathbb{N}$ be the function that outputs the length of its input.

Lemma 6.1. *The following diagram commutes:*

$$\begin{array}{ccc}
 list(\mathbb{N}) & \xrightarrow{len} & \mathbb{N} \\
 qck_{list(\mathbb{N})} \downarrow & & \downarrow qck_{\mathbb{N}} \\
 Flist(\mathbb{N}) & \xrightarrow{F(len)} & F\mathbb{N}
 \end{array}$$

PROOF. This can be proved by strong induction on the length of the input list. If the list is empty, then its length is 0 and the diagram commutes. If, however, the list not not empty, there will be a bijection between the recursive calls to lists of size n' to sampling n' in $qck_{\mathbb{N}}$. By using the strong inductive hypothesis to these lists of size n' , we can conclude. \square

$qck_{\mathbb{N}} = \text{fix } f : \mathbb{N} \rightarrow F\mathbb{N}. \lambda n : \mathbb{N}.$

if n then

produce 0

else

charge $n - 1$; ⁽¹⁾

$x \leftarrow \text{rand } n$; ⁽²⁾

(force f) x ; ⁽³⁾

(force f) $(n - x)$; ⁽⁴⁾

produce n ⁽⁵⁾

$\llbracket qck_{\mathbb{N}} \rrbracket_{EC} = \bigsqcup_n F^n(\perp)$, where

$F = \lambda T. \lambda n. \text{if } n \text{ then}$

0

else

$(n - 1)$ ⁽¹⁾ + $\sum_i \frac{1}{n}$ ⁽²⁾ ($T(i)$) ⁽³⁾ + $T(n - i)$ ⁽⁴⁾ + 0 ⁽⁵⁾

Fig. 9. Quicksort over natural numbers and its denotational expected cost

We can now conclude our analysis, since by the soundness theorem and the commutative diagram above, the expected cost of quickSort is given by $\pi_1 \circ qck_{\mathbb{N}} \circ \text{len}$. By unfolding the definition of $\pi_1 \circ \llbracket qck_{\mathbb{N}} \rrbracket_{EC}$, we obtain the following expression $\text{fix } (f \mapsto n \mapsto \text{ifZero } n \text{ then } 1 \text{ else } (n - 1) + \sum_i \frac{1}{n} (f(i) + f(n - i)))$ which can be further simplified to the recurrence relation:

$$T(0) = 0$$

$$T(n) = n - 1 + \frac{2}{n} \sum_{i=0}^{i-1} T(i)$$

This allows us to conclude that quickSort has an expected cost of $O(n \log(n))$.

6.3 Random Walks

For this example we are interested in the symmetric random walk over the natural numbers. At every point n the probability of moving to $n - 1$ or $n + 1$ is $\frac{1}{2}$. Furthermore, we are assuming the variant where at 0 you move to 1 with probability 1. We can write a program that simulates such a random walk with a point of departure $i : \mathbb{N}$ and a point of arrival $j : \mathbb{N}$:

$\text{randomWalk} = \mu f : \mathbb{N} \rightarrow \mathbb{N} \rightarrow F1. \lambda i : \mathbb{N} j : \mathbb{N}.$

if $i = j$ then

produce ()

else

charge 1;

if i then

(force f) 1 j

else

((force f) $(i - 1) j$) \oplus ((force f) $(i + 1) j$)

The program receives the starting and end points, i and j , respectively, as arguments, and if they are equal, you stop the random walk. Otherwise, you take one step of the random walk, i.e. you take step to either $i - 1$ or $i + 1$ with equal probability, with the exception of when $i = 0$, in which

case you go to 1. This iterative behaviour can be straightforwardly captured with recursion, as illustrated by the program above.

It is now possible to compute the expected value on the number of rounds that are necessary in order to reach your target. This is given by the following two-argument recursive relation.

$$\begin{aligned} T(i, i) &= 0 \\ T(0, j) &= 1 + T(1, j) \\ T(i, j) &= 1 + \frac{1}{2}(T(i - 1, j) + T(i + 1, j)) \end{aligned}$$

This recurrence relation is well-known in the theory of Markov chains – see [34] for an introduction. Something interesting about it is that when $i > j$, this stochastic process reduces to the symmetric random walk without an absorbing state, which is known to have ∞ expected cost.

7 RELATED WORK

There has been much work done on logic and language techniques for reasoning about the cost of programs.

Type Theories for Cost Analysis. Recent work [13, 33] have developed (in)equational theories for reasoning about costs of programs inside a modal dependently-typed CBPV metalanguage. Their framework can reason about monadic effects by using the writer monad transformer, similarly to `cert`'s cost semantics, but, due to being inside a total dependent type theory, it can only represent total programs and finitely supported distributions— i.e. it cannot represent the geometric distribution. Furthermore, it is unclear how one would go about extending their framework with continuous distributions.

Other work has focused in designing type theories for doing relational reasoning of programs [7, 35]. Even though these approaches can reason about functional programs as well, they are limited to deterministic programs.

In work by Avanzini et al. [2], the authors define a graded, substructural type system for reasoning about expected cost of functional programs, even using a randomized quicksort as an example. One of the main limitations of their system with respect to `cert` is that, due to the substructural invariants of their type system, it can only type check a limited subset of the programs that `cert` can. For instance, it cannot typecheck common functional idioms like `fold` and `map` functions over lists. Furthermore, they have not addressed how feasible type checking in their system is or if it is even decidable, which in the context of type-based reasoning is an important property to have.

In other work by Avanzini et al. [1], the authors describe a continuation passing style (CPS) transformation into a metalanguage for reasoning about expected cost of programs. Compared to `cert`, both metalanguages can handle functional programming, though their language is restricted to a CBV semantics. Furthermore, using continuation passing style to reason about programs creates undesired gaps between the transformed and original programs which are avoided when using the expected cost monad's direct-style reasoning.

Automatic Resource Analysis. One fruitful research direction has been the automatic amortized resource analysis (AARA) [16, 17, 32] which uses a type system to annotate programs with their cost and automatically infer the cost of the program. By now, these techniques have been extended to reason about recursive types [14], probabilistic programs [38] and programs with local state [28].

Something quite appealing about their approach is that it is completely automatic, whereas our approach requires solving a, possibly hard, recurrence relation by hand. That being said, their system can only accommodate polynomial bounds, meaning that they cannot infer the $n \log(n)$

bound for the probabilistic quicksort like we do. Recently, AARA has been extended to accommodate exponential bounds [18] in deterministic programs, though it is still unclear if the same technique can be extended to the probabilistic setting, meaning that they cannot analyse the behaviour of exponentially slow programs such as the expected coin tosses one.

There have been other type-based approach to automatically reasoning about cost of programs, such as the language TiML [39]. This language allows users to annotate type signatures with cost-bounds and the type checking algorithm will infer and check these bounds. The main limitation of TiML in comparison to our work is that it cannot handle probabilistic programs. There has also been work done on automated reasoning about cost for first-order probabilistic programs by Avanzini et al. [3]. The main limitation of this work when compared to `cert` is that it can only handle first-order imperative programs.

It is an interesting line of future work understanding to what extent solving recurrence relations can be automated in the context of `cert`.

Recurrence for Expected Cost. There has been some work done in exploring languages for expressing recurrence relations for expected cost. For example, [36] provides a language for representing probabilistic recurrence relations and a tool for analysing their tail-bounds. The main drawback of these approaches is that the languages are not very expressive. In particular they do not have higher-order functions.

Leutgeb et al. [26] define a first-order probabilistic functional language for manipulating data structures and automatically infer bounds on the expected cost of programs. The main limitation of their approach compared to ours is that their language is first-order.

Reasoning about expected cost has also been explored for imperative languages. For instance, in [4] the authors develop a weakest pre-condition calculus for reasoning about the expected cost of programs. Again, they can only reason about first-order imperative programs.

8 CONCLUSION AND FUTURE WORK

In this work we have presented `cert`, a metalanguage for reasoning about expected cost of recursive probabilistic programs. It extends the existing work of [24] to the probabilistic setting. We have proposed two different semantics, one based on the writer monad transformer while the second one uses a novel *expected cost* monad. Furthermore, we have showed that in the absence of unbounded recursion, these two semantics coincide, while when programming with subprobability distributions we have proved that the expected cost semantics is an upper bound to the cost semantics.

We have justified the versatility of our expected cost semantics by presenting a few case-studies. In particular, the expected cost semantics obtains, compositionally, the familiar recurrence cost relations for non-trivial programs. In particular, for the randomized quicksort algorithm, the semantic recurrence relation recovers the $O(n \log n)$ bound.

We conjecture that the techniques presented in this paper can be extended in various other tantalizing directions. For instance, it would be interesting to generalize the construction of the expected cost monad so that it can reason about other kinds of effects. For instance, when reasoning about non-deterministic programs it is useful to give bounds on the worse/best cost.

REFERENCES

- [1] Martin Avanzini, Gilles Barthe, and Ugo Dal Lago. 2021. On continuation-passing transformations and expected cost analysis. In *International Conference on Functional Programming (ICFP)*.
- [2] Martin Avanzini, Ugo Dal Lago, and Alexis Ghyselen. 2019. Type-based complexity analysis of probabilistic functional programs. In *Logic in Computer Science (LICS)*.
- [3] Martin Avanzini, Georg Moser, and Michael Schaper. 2020. A modular cost analysis for probabilistic programs. In *Object-oriented Programming, Systems, Languages, and Applications (OOPSLA)*.

- [4] Kevin Batz, Benjamin Lucien Kaminski, Joost-Pieter Katoen, Christoph Matheja, and Lena Verscht. 2023. A calculus for amortized expected runtimes.
- [5] Francis Borceux. 1994. *Handbook of categorical algebra: volume 1, Basic category theory*. Vol. 1. Cambridge University Press.
- [6] Krishnendu Chatterjee, Hongfei Fu, Petr Novotný, and Rouzbeh Hasheminezhad. 2016. Algorithmic analysis of qualitative and quantitative termination problems for affine probabilistic programs. In *Principles of Programming Languages (POPL)*.
- [7] Ezgi Çiçek, Gilles Barthe, Marco Gaboardi, Deepak Garg, and Jan Hoffmann. 2017. Relational cost analysis. In *Principles of Programming Languages (POPL)*.
- [8] Youyou Cong, Leo Osvald, Grégory M Essertel, and Tiark Rompf. 2019. Compiling with Continuations, or without? Whatever. *Proceedings of the ACM on Programming Languages* 3, ICFP (2019), 1–28.
- [9] Joseph W Cutler, Daniel R Licata, and Norman Danner. 2020. Denotational recurrence extraction for amortized analysis.
- [10] Norman Danner, Daniel R Licata, and Ramyaa Ramyaa. 2015. Denotational cost semantics for functional languages with inductive types. In *International Conference on Functional Programming (ICFP)*.
- [11] Norman Danner, Jennifer Paykin, and James S Royer. 2013. A static cost analysis for a higher-order language. In *7th workshop on Programming languages meets program verification*.
- [12] Thomas Ehrhard, Michele Pagani, and Christine Tasson. 2017. Measurable cones and stable, measurable functions: a model for probabilistic higher-order programming. *Proceedings of the ACM on Programming Languages* POPL (2017).
- [13] Harrison Grodin, Yue Niu, Jonathan Sterling, and Robert Harper. 2023. Decalf: A Directed, Effectful Cost-Aware Logical Framework. In *Principles of Programming Languages (POPL)*.
- [14] J. Grosen, D. M. Kahn, and J. Hoffmann. 2023. Automatic Amortized Resource Analysis with Regular Recursive Types. In *Logic in Computer Science (LICS)*.
- [15] Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang. 2017. A convenient category for higher-order probability theory. In *Logic in Computer Science (LICS)*.
- [16] Jan Hoffmann, Ankush Das, and Shu-Chun Weng. 2017. Towards automatic resource bound analysis for OCaml. In *Symposium on Principles of Programming Languages (POPL)*.
- [17] Jan Hoffmann and Zhong Shao. 2015. Automatic static cost analysis for parallel programs. In *European Symposium on Programming (ESOP)*.
- [18] David M Kahn and Jan Hoffmann. 2020. Exponential automatic amortized resource analysis. In *Foundations of Software Science and Computation Structures (FoSSaCS)*.
- [19] Benjamin Lucien Kaminski, Joost-Pieter Katoen, Christoph Matheja, and Federico Olmedo. 2016. Weakest precondition reasoning for expected run-times of probabilistic programs. In *European Symposium on Programming (ESOP)*.
- [20] Ohad Kammar and Dylan McDermott. 2018. Factorisation systems for logical relations and monadic lifting in type-and-effect system semantics. *Electronic notes in theoretical computer science* (2018).
- [21] Richard M Karp. 1994. Probabilistic recurrence relations. *Journal of the ACM (JACM)* (1994).
- [22] Shin-ya Katsumata. 2005. A semantic formulation of $\top\top$ -lifting and logical predicates for computational metalanguage. In *International Workshop on Computer Science Logic*. Springer, 87–102.
- [23] Shin-ya Katsumata. 2013. Relating computational effects by $\top\top$ -lifting. *Information and Computation* (2013).
- [24] GA Kavvos, Edward Morehouse, Daniel R Licata, and Norman Danner. 2019. Recurrence extraction for functional programs through call-by-push-value. In *Principles of Programming Languages (POPL)*.
- [25] Satoshi Kura, Natsuki Urabe, and Ichiro Hasuo. 2019. Tail probabilities for randomized program runtimes via martingales for higher moments. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*.
- [26] Lorenz Leutgeb, Georg Moser, and Florian Zuleger. 2022. Automated expected amortised cost analysis of probabilistic data structures. In *Computer Aided Verification (CAV)*.
- [27] Paul Blain Levy. 2001. *Call-by-push-value*. Ph. D. Dissertation.
- [28] Benjamin Lichtman and Jan Hoffmann. 2017. Arrays and references in resource aware ML. In *Formal Structures for Computation and Deduction (FSCD 2017)*.
- [29] Luke Maurer, Paul Downen, Zena M Ariola, and Simon Peyton Jones. 2017. Compiling without continuations. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*. 482–494.
- [30] E Moggi. 1989. Computational lambda-calculus and monads. In *Logic in Computer Science (LICS)*.
- [31] Rajeev Motwani and Prabhakar Raghavan. 1995. *Randomized algorithms*. Cambridge university press.
- [32] Van Chan Ngo, Quentin Carbonneaux, and Jan Hoffmann. 2018. Bounded expectations: resource analysis for probabilistic programs. In *Programming Language Design and Implementation (PLDI)*.
- [33] Yue Niu, Jonathan Sterling, Harrison Grodin, and Robert Harper. 2022. A cost-aware logical framework. In *Principles of Programming Languages (POPL)*.
- [34] James R Norris. 1998. *Markov chains*. Number 2. Cambridge university press.

- [35] Vineet Rajani, Marco Gaboardi, Deepak Garg, and Jan Hoffmann. 2021. A unifying type-theory for higher-order (amortized) cost analysis. In *Symposium on Principles of Programming Languages (POPL)*.
- [36] Yican Sun, Hongfei Fu, Krishnendu Chatterjee, and Amir Kafshdar Goharshady. 2023. Automated Tail Bound Analysis for Probabilistic Recurrence Relations. In *Computer Aided Verification (CAV)*.
- [37] Matthijs Vákár, Ohad Kammar, and Sam Staton. 2019. A domain theory for statistical probabilistic programming. In *Principles of Programming Languages (POPL)*.
- [38] Di Wang, David M Kahn, and Jan Hoffmann. 2020. Raising expectations: automating expected cost analysis with types. *International Conference on Functional Programming (ICFP)*.
- [39] Peng Wang, Di Wang, and Adam Chlipala. 2017. TiML: a functional language for practical complexity analysis with invariants. In *Object-oriented Programming, Systems, Languages, and Applications (OOPSLA)*.

A MONADIC SEMANTICS OF CBPV

Let \mathcal{C} be a Cartesian closed category and $T : \mathcal{C} \rightarrow \mathcal{C}$ a strong monad over it. An alternative definition of monads is it being a triple (T, η, μ) , where T and η are natural transformations as before, but $\mu : T^2 \rightarrow T$, the multiplication, replaces the bind natural transformation. The monad laws under this definition become:

$$\begin{array}{ccc} T^3 & \xrightarrow{T\mu} & T^2 \\ \mu_T \downarrow & & \downarrow \mu \\ T^2 & \xrightarrow{\mu} & T \end{array} \qquad \begin{array}{ccccc} T & \xrightarrow{T\eta} & T^2 & \xleftarrow{\eta_T} & T \\ & \searrow 1 & \downarrow \mu & \swarrow 1 & \\ & & T & & \end{array}$$

It is possible to show that these definitions are equivalent: given $\text{bind } (-)^\#$, the multiplication can be defined as $\mu = \text{id}_{TA}^\#$. Conversely, given a multiplication, the bind is defined as $f^\# = Tf; \mu$. This alternative definition is a bit better suited for the original purposes of monads, where it was used as a unifying way of representing concepts from universal algebra.

This alternative presentation lend itself quite well to the semantics of CBPV-based calculi where, given a monad T , computation types denote T -algebras:

Definition A.1. A T -algebra is a pair (A, α) , where A is a \mathcal{C} object and $\alpha : TA \rightarrow A$ is a morphism, such that

$$\begin{array}{ccc} A & \xrightarrow{\eta_A} & TA \\ & \searrow \text{id}_A & \downarrow \alpha \\ & & A \end{array} \qquad \begin{array}{ccc} T^2A & \xrightarrow{\mu_A} & TA \\ T\alpha \downarrow & & \downarrow \alpha \\ TA & \xrightarrow{\alpha} & A \end{array}$$

Given a T -algebra (A, α) we denote by A_\bullet the object of the T -algebra.

Example A.2. Given an object A , the pair (TA, μ_A) is a T -algebra, where the algebra axioms follow from the monad laws.

Example A.3. Given a T -algebra (A, α) and an objects B , we can equip $B \rightarrow B$ with the T -algebra structure $\alpha_{B \rightarrow A} = \varepsilon; B \Rightarrow (st; T(ev; \alpha))$, where $\varepsilon_A : A \rightarrow (B \Rightarrow (B \times A))$ is the unit of the Cartesian closed adjunction.

Algebras and their morphisms can be organized as a category, frequently denoted by \mathcal{C}^T . However, for the purposes of CBPV a different category is used:

Definition A.4. The category $\tilde{\mathcal{C}}^T$ is the full subcategory of \mathcal{C} that contain T -algebras as objects. This category is also called the category of algebras and plain maps.

The idea is that values are interpreted as objects in \mathcal{C} while computation types are T -algebras. Assuming the only the base type in the calculus to be \mathbb{N} and an object \mathbb{N} in the base category, The interpretation of values and computations are as follows:

$$\begin{aligned} \llbracket \mathbb{N} \rrbracket^v &= \mathbb{N} \\ \llbracket U\bar{\tau} \rrbracket^v &= \llbracket \bar{\tau} \rrbracket^c \\ \llbracket \tau_1 \times \tau_2 \rrbracket^v &= \llbracket \tau_1 \rrbracket^v \times \llbracket \tau_2 \rrbracket^v \\ \llbracket F\tau \rrbracket^c &= (T \llbracket \tau \rrbracket^v, \mu_{\llbracket \tau \rrbracket^v}) \\ \llbracket \tau \rightarrow \bar{\tau} \rrbracket^c &= (\llbracket \tau \rrbracket^v \Rightarrow \llbracket \bar{\tau} \rrbracket^c, \alpha_{\llbracket \tau \rrbracket^v \rightarrow \llbracket \bar{\tau} \rrbracket^c}) \end{aligned}$$

$$\begin{array}{c}
\text{VAR} \\
\hline
\Gamma_1 \times (\tau \times \Gamma_2) \xrightarrow{! \times \tau_1} \tau
\end{array}
\quad
\begin{array}{c}
\text{IF} \\
\hline
\frac{\Gamma \xrightarrow{V} \mathbb{N} \quad \Gamma \xrightarrow{t} \bar{\tau} \quad \Gamma \xrightarrow{u} \bar{\tau}}{\Gamma \xrightarrow{\langle id, V \rangle; [t, (! \mathbb{N}; u)]} \bar{\tau}}
\end{array}
\quad
\begin{array}{c}
\text{ABSTRACTION} \\
\hline
\frac{\Gamma \times \tau \xrightarrow{t} \bar{\tau}}{\Gamma \xrightarrow{\Lambda_\Gamma; \tau \Rightarrow t} \tau \Rightarrow \bar{\tau}}
\end{array}
\quad
\begin{array}{c}
\text{APPLICATION} \\
\hline
\frac{\Gamma \xrightarrow{V} \tau \quad \Gamma \xrightarrow{t}: \tau \rightarrow \bar{\tau}}{\Gamma \xrightarrow{\langle t, V \rangle; ev} \bar{\tau}}
\end{array}$$

$$\begin{array}{c}
\text{PRODUCE} \\
\hline
\frac{\Gamma \xrightarrow{V} \tau}{\Gamma \xrightarrow{V; \eta_\tau} T\tau}
\end{array}
\quad
\begin{array}{c}
\text{BIND} \\
\hline
\frac{\Gamma \xrightarrow{t} T\tau' \quad \Gamma \times \tau' \xrightarrow{u} (\bar{\tau}, \alpha_{\bar{\tau}})}{\Gamma \xrightarrow{\langle id_\Gamma, t \rangle; st; Tu; \alpha_{\bar{\tau}}} (\bar{\tau}, \alpha_{\bar{\tau}})}
\end{array}
\quad
\begin{array}{c}
\text{THUNK} \\
\hline
\frac{\Gamma \xrightarrow{t} (\bar{\tau}, \alpha_{\bar{\tau}})}{\Gamma \xrightarrow{t} \bar{\tau}}
\end{array}
\quad
\begin{array}{c}
\text{FORCE} \\
\hline
\frac{\Gamma \xrightarrow{V} \bar{\tau}}{\Gamma \xrightarrow{V} (\bar{\tau}, \alpha_{\bar{\tau}})}
\end{array}$$

$$\begin{array}{c}
\text{LET} \\
\hline
\frac{\Gamma \xrightarrow{V} \tau' \quad \Gamma \times \tau' \xrightarrow{t} \bar{\tau}}{\Gamma \xrightarrow{\langle id_\Gamma, V \rangle; t} \bar{\tau}}
\end{array}
\quad
\begin{array}{c}
\text{PAIR} \\
\hline
\frac{\Gamma \xrightarrow{t_1} \tau_1 \quad \Gamma \xrightarrow{t_2} \tau_2}{\Gamma \xrightarrow{\langle t_1, t_2 \rangle} \tau_1 \times \tau_2}
\end{array}
\quad
\begin{array}{c}
\text{UNPAIR} \\
\hline
\frac{\Gamma \xrightarrow{V} \tau_1 \times \tau_2 \quad \Gamma \times \tau_1 \times \tau_2 \xrightarrow{t} \bar{\tau}}{\Gamma \xrightarrow{\langle id, V \rangle; t} \bar{\tau}}
\end{array}$$

Fig. 10. CBPV monadic semantics

It is also possible to give semantics to the terms of the language as depicted in Figure 10. The semantics of if-statements use the fact that $\mathbb{N} \cong 1 + \mathbb{N}$, so you can define its semantics by using the universal property of coproducts $[t, (!\mathbb{N}; u)]$, where $!_A : A \rightarrow 1$ is the unique arrow into the terminal object. The abstraction and application rule use the adjoint structure (Λ, ev) , of Cartesian closed categories, where Λ and ev are the unit and counit of the adjunction, respectively. The produce rule uses the unit of the monad while the bind rule is the sequential composition of a free algebra with a non-free algebra and, therefore, requires applying the functor T and using the algebra structure of the output – when the output map is a free algebra, this operation is equal to the bind of the monad. Thunk and force are basically no-ops in this semantics, while the rules let and unpair are sequential compositions. Pair is the universal property of products.

A.1 Equational presentation of cert

For the sake of simplicity of the equational theory, we will assume the barycentric operations \oplus_p .

In Figure 11 we present the non-structural equations of **cert**. The left-hand side is present in every CBPV calculus with natural numbers and recursion, where the recursion equation is the last one. The right-hand side is split in two blocks: the first block are the barycentric algebra equations, the second one are the monoid equations and the last one are the list equations.

B DENOTATIONAL SOUNDNESS PROOF

Lemma B.1. *If $(r, \nu) \mathcal{C}_{F\tau} \mu$ then for every pair of functions $f_1 : \llbracket \tau \rrbracket_{EC} \rightarrow \mathbb{R}$ and $f_2 : \llbracket \tau \rrbracket_{CS} \rightarrow \mathbb{R}$, such that for every $a_1 \mathcal{V}_\tau a_2$, $f_1(a_1) \leq f_2(a_2)$, $\int f_1 d\nu \leq \int f_2 d\mu_2$, where μ_2 is the second marginal of μ .*

PROOF. Since by assumption $(r, \nu) \mathcal{C}_{F\tau} \mu$, there is a coupling γ over the support of \mathcal{V}_τ , which allows us to conclude:

$$\int f_1 d\nu \leq \int \frac{1}{2}(f_1 + f_2) d\gamma \leq \int f_2 d\mu_2$$

The equalities above hold because, in the support of γ , $f_1(a_1) \leq f_2(a_2)$, making $f_1 \leq \frac{1}{2}(f_1 + f_2) \leq f_2$ and since γ is a joint distribution with marginals ν and μ_2 , we have the (in)equality of integrals above. \square

The following lemma is the most technical aspect of the soundness proof and, intuitively, is saying that the logical relations for computation types can be equipped with "algebra" structures.

ifZero 0 then t else $u \equiv t$	$t \oplus_0 u \equiv t$
ifZero $(n + 1)$ then t else $u \equiv u$	$t \oplus_p u \equiv u \oplus_{1-p} t$
$t \equiv \text{ifZero } x \text{ then } t \text{ else } t$	$t \oplus_p t \equiv t$
$(\lambda x. t) V \equiv t\{V/x\}$	$t \oplus_p (u \oplus_q t') \equiv (t \oplus_{pq} u) \oplus_{\frac{p(1-q)}{1-pq}} t'$
let x be V in $t \equiv t\{V/x\}$	charge n ; charge $m \equiv \text{charge } n + m$
$t \equiv \lambda x. t x$	charge 0; $t \equiv t$
$x \leftarrow t; (\lambda y. u) \equiv \lambda y. (x \leftarrow t; u)$	t ; charge 0 $\equiv t$
force (think t) $\equiv t$	
think (force V) $\equiv V$	
$x \leftarrow (\text{produce } V); t \equiv t\{V/x\}$	case nil of nil $\Rightarrow t \mid \text{cons } x \text{ xs} \Rightarrow u \equiv t$
$x \leftarrow t; \text{produce } x \equiv t$	case (cons $V_1 V_2$) of nil $\Rightarrow t \mid \text{cons } x \text{ xs} \Rightarrow u \equiv u\{V_1, V_2/x, \text{xs}\}$
fix $x. t = t\{(\text{think } (\text{fix } x. t)) / x\}$	$t \equiv \text{case } y \text{ of nil} \Rightarrow t\{\text{nil}/y\} \mid \text{cons } x \text{ xs} \Rightarrow t\{\text{cons } x \text{ xs}/y\}$

Fig. 11. cert equational theory

Furthermore, since we are proving two similar looking theorems for the probabilistic and subprobabilistic cases, and the soundness proof in both cases is basically the same, we will only present the proof to the subprobabilistic case, and highlight in the proof what would differ for the probabilistic case.

Lemma B.2. *Let τ_1, τ_2 and $\bar{\tau}$ be types and $f_1 : \llbracket \tau_1 \rrbracket_{CS}^v \times \llbracket \tau_2 \rrbracket_{CS}^v \rightarrow \llbracket \bar{\tau} \rrbracket_{CS}^c$, and $f_2 : \llbracket \tau_1 \rrbracket_{EC}^v \times \llbracket \tau_2 \rrbracket_{EC}^v \rightarrow \llbracket \bar{\tau} \rrbracket_{EC}^c$ be ω Qbs morphisms such that $f_1 \times f_2$, when the input is restricted to $\mathcal{V}_{\tau_1} \times \mathcal{V}_{\tau_2}$, the output is restricts to $C_{\bar{\tau}}$. It is true that $(st; T_1(f_1); \alpha_{\bar{\tau}}) \times (st; T_2(f_2); \alpha_{\bar{\tau}})$, when its input is restricted to $\mathcal{V}_{\tau_1} \times C_{F_{\tau_2}}$, has its output still be restricted to $C_{\bar{\tau}}$.*

PROOF. This can be proved by induction on the computation type $\bar{\tau}$:

$F\tau$: In order to prove $f_1^\#(r, \nu) C_{F\tau} f_2^\#(\mu)$ we have to prove that their expected costs are related by the inequality given by the definition of $C_{F\tau}$ and show that there is a coupling over ν and μ_2 , where μ_2 is the second marginal of μ , such that it factors through the inclusion $P_{\leq 1}(\mathcal{V}_{\tau'}) \hookrightarrow P_{\leq 1}(\llbracket \tau' \rrbracket_{CS}^v \times \llbracket \tau' \rrbracket_{EC}^v)$.

By unfolding the definitions, we get

$$\begin{aligned} \pi_1(f_1^\#(r, \nu)) &= r + \int (\pi_1 \circ f_1) d\nu \\ \mathbb{E}(f_2^\#(\mu)) &= \int \int n \|f_2(a)\| \mu(dn, da) + \int n d(f_2^\#(\mu)_1) \end{aligned}$$

In the second expression, the left hand side term being summed corresponds to the expected cost of the input while the second one corresponds to the cost of the continuation. As such, it is sensible that, in order to reason about their difference, we should reason individually about $r - \int \int n \|f_2(a)\|$ and $\int (\pi_1 \circ f_1) d\nu - \int n d(f_2^\#(\mu))$, and both should be greater than 0. The first inequality is immediate:

$$\int \int n \|f_2(a)\| d\mu \leq \int \int n d\mu \leq r$$

For the second expression, assuming $\forall a' \mathcal{V}_{\tau'} a, \mathbb{E}(f_2(a)) \leq (\pi_1 \circ f_1)(a')$, we can apply Lemma B.1 and use the equality $\int n \, d(f_2^\#(\mu)_1) = \int \mathbb{E}(f_2(a)_1) \, d\mu_2$.

By adding these two inequalities we obtain exactly the first condition of the relation $C_{F\tau}$. In the probabilistic case every inequality is an equality, since $\|f(a)\| = 1$ and the inequalities in the definition of $C_{F\tau}$ would be equalities as well. The second condition follows from observing that when restricting the domain of $f_1 \times f_2$ to \mathcal{V}_τ , we can extract from it a morphism $g : \mathcal{V}_\tau \rightarrow P_{\leq 1}(\mathcal{V}_{\tau'})$ such that, given inputs (v_1, v_2) , the marginals of $g(v_1, v_2)$ are equal to $\pi_2(f_1(v_1))$ and $f_2(v_2)_2$ since, by assumption, $f_1(v_1) C_{F\tau'} f_2(v_2)$.

Given this function, we define the coupling $g^\#(\mu')$, where μ' is the coupling given by the “witness” of $(r, v) C_{F\tau} \mu$. Showing that it has the right marginals follows from linearity of the marginal function, concluding the proof. This part of the proof remains the same in the probabilistic case

$\tau \rightarrow \bar{\tau}$: This case relies more on notation and, therefore, in order to simplify the presentation, we will rely on the symmetry of f_1 and f_2 and work on the generic expression $st; T(f); \alpha_{\tau \rightarrow \bar{\tau}}$ that can be instantiated to both f_1 and f_2 .

By definition of $C_{\tau \rightarrow \bar{\tau}}$, in order to define a morphism $\mathcal{V}_{\tau_1} \times \mathcal{V}_{\tau_2} \rightarrow C_{\tau \rightarrow \bar{\tau}}$, it suffices to define its transpose $\mathcal{V}_\tau \times (\mathcal{V}_{\tau_1} \times \mathcal{V}_{\tau_2}) \rightarrow C_{\bar{\tau}}$. Since the algebra structure of $\alpha_{\tau \rightarrow \bar{\tau}}$ is defined as $\eta; id_\tau \Rightarrow (st; T(ev); \alpha_{\bar{\tau}})$, we want to show that the map $id_\tau \times (st; Tf; \eta; id_\tau \Rightarrow (st; T(ev); \alpha_{\bar{\tau}})); ev$, i.e. can be rewritten in the format $st; T(f'); \alpha_{\bar{\tau}}$, so that we can apply the induction hypothesis. This equation holds, up to isomorphism, by the following commutative diagram:

$$\begin{array}{ccccccc}
\tau \times (\tau_1 \times T\tau_2) & \xrightarrow{id_\tau \times st} & \tau \times T(\tau_1 \times \tau_2) & \xrightarrow{\tau \times Tf} & \tau \times T(\tau \Rightarrow \bar{\tau}) & \xrightarrow{\tau \times \eta} & \tau \times (\tau \Rightarrow (\tau \times T(\tau \Rightarrow \bar{\tau}))) \\
\downarrow a & & \downarrow st & & \parallel & \swarrow ev & \downarrow id_\tau \times (id_\tau \Rightarrow (st; T(ev); \alpha_{\bar{\tau}})) \\
& & & & \tau \times T(\tau \Rightarrow \bar{\tau}) & & \tau \times (\tau \Rightarrow \bar{\tau}) \\
& & & & \downarrow st & & \downarrow ev \\
(\tau \times \tau_1) \times T\tau_2 & \xrightarrow{st; T(a^{-1})} & T(\tau \times (\tau_1 \times \tau_2)) & \xrightarrow{T(\tau \times f)} & T(\tau \times (\tau \Rightarrow \bar{\tau})) & \xrightarrow{Tev} & T\bar{\tau} \xrightarrow{\alpha_{\bar{\tau}}} \bar{\tau}
\end{array}$$

From left to right, the first diagram commutes by definition of strong monad, the second commutes from naturality of the strength of T , the triangular diagram commutes by the Cartesian closed adjunction and the final diagram commutes by naturality of ev . \square

We are interested in the case where the type τ_1 will be a context Γ . We can now prove the full statement.

PROOF. The proof follows from mutual induction on $\Gamma \vdash^o V : \tau$ and $\Gamma \vdash^c t : \bar{\tau}$. Many of the cases follow by just applying the induction hypothesis or by assumptions of theorem. We go over the most interesting cases:

Comp This case follows from Lemma B.2.

Fix This theorem follows from the induction hypothesis and from the fact that the relations $C_{\bar{\tau}}^c$ are closed under suprema of ascending chains.

Produce First apply the induction hypothesis to V and assume that $\llbracket V \rrbracket_1 = v_1$ and $\llbracket V \rrbracket_2 = v_2$.
By construction, $\eta^{T_1}(v_1) C_{F\tau}^v \eta^{T_2}(v_2)$, since they both have the same expected value and the coupling is $\delta_{(v_1, v_2)}$.

Case By applying the inductive hypothesis to V we may do case analysis on it and if it is the empty list, we use the inductive hypothesis on t and, otherwise, we use the inductive hypothesis on u .

□

C OPERATIONAL SOUNDNESS PROOF

The soundness proof hinges on the the following lemma that can be proved by induction on computation types.

Lemma C.1. *If $\Gamma \vdash_c t : F\tau$ and $x : \tau, \Gamma \vdash_c u : \bar{\tau}$ then $\llbracket \Downarrow (x \leftarrow t; u) \rrbracket = (v \mapsto \llbracket \Downarrow (u\{v/x\}) \rrbracket)^\#(\alpha(\llbracket \Downarrow (t) \rrbracket))$*

We now prove the soundness theorem.

PROOF. Proof by induction on n and on t . The base case $n = 0$ is trivial because \perp is the least element and such an element is preserved by the algebra structure. We present the base term cases and the sequencing case. The inductive ones follow from the inductive hypothesis, Lemma C.1 and the monotonicity of the algebra structure α .

Terminal : The evaluation rules for terminal computation output the unit of the monad, which allows us to conclude $(\alpha \circ \eta)(a) = a \leq \llbracket t \rrbracket_{CS}$.

Charge : $\alpha(\delta_{(r, \delta_{(0, ())})}) = \delta_{(r, ())} \leq \llbracket \text{charge } r \rrbracket_{CS}$.

Sampling : $\alpha(\delta_{(0, \delta_0 \otimes \lambda)}) = \delta_0 \otimes \lambda \leq \llbracket \text{uniform} \rrbracket_{CS}$.

Seq : We illustrate an inductive case. The rest follow a similar pattern.

$$\begin{aligned} \llbracket \Downarrow (x \leftarrow t; u) \rrbracket &= \\ (\alpha \circ (v \mapsto \llbracket \Downarrow (u\{v/x\}) \rrbracket)^\#) &(\alpha(\llbracket \Downarrow (t) \rrbracket)) \leq \\ (\alpha \circ (v \mapsto \llbracket \Downarrow (u\{v/x\}) \rrbracket)^\#) &(\llbracket t \rrbracket_{CS}) \leq \\ \llbracket x \leftarrow t; u \rrbracket_{CS} & \end{aligned}$$

where the last inequality follows from monotonicity of the algebra structure and the algebra axioms.

□

D OPERATIONAL ADEQUACY PROOF

In order to prove the fundamental theorem of logical relations we require a couple of lemmas that are proved by induction.

Lemma D.1. *If $t \triangleleft f$ (resp. $V \triangleright v$) then $\llbracket t \rrbracket_{CS} \leq f$ (resp. $\llbracket V \rrbracket_{CS} \leq v$).*

Lemma D.2. *The relations \triangleleft are closed under ascending chains.*

Lemma D.3. *If $t \triangleleft f$ and $t \equiv u$, then $u \triangleleft f$.*

PROOF. The proof follows by induction on the computation type and the equational theory. □

We need a lemma similar to Lemma B.2 which once again follows by induction on the computation type.

Lemma D.4. *Let Γ be a context, τ_1 and $\bar{\tau}$ types, $\Gamma, x : \tau_1 \vdash t : \bar{\tau}$ a computation and $f : \llbracket \Gamma \rrbracket_{CS}^v \times \llbracket \tau_1 \rrbracket_{CS}^v \rightarrow \llbracket \bar{\tau} \rrbracket_{CS}^c$ a ω Qbs morphism such that for every $\cdot \vdash_v V : \tau_1$ with $V \triangleright v$, $t\{G, V/\Gamma, x\} \triangleleft f(\gamma, v)$, then $(x \leftarrow u; t) \triangleleft (\alpha \circ Tf(\gamma))(v)$, whenever $u \triangleleft v$.*

Theorem D.5 (Fundamental Theorem of Logical Relations). *If $\Gamma \vdash_c t : \bar{\tau}$ (resp. $\Gamma \vdash_v V : \tau$) and $V_i \triangleright v_i$ then $t\{V_1, \dots, V_n/x_1, \dots, x_n\} \llcorner \llbracket t \rrbracket_{CS}(v_1, \dots, v_n)$ (resp. $V\{V_1, \dots, V_n/x_1, \dots, x_n\} \triangleright \llbracket V \rrbracket_{CS}(v_1, \dots, v_n)$).*

PROOF. The proof follows by mutual induction on the typing derivations of t and V .

Var : $x_i\{V_i/x_i\} = V_i$ which implies the conclusion by the assumption $V_i \triangleright v_i$.

Arithmetic constants : Follows by inspection.

Pair : Follows directly from the induction hypothesis.

Charge : By unfolding the definitions, $\llbracket \Downarrow \text{charge } V \rrbracket_{CS} = \delta_{\llbracket V \rrbracket_*} = \llbracket \text{charge } V \rrbracket_{CS}$

Sample : By unfolding the definitions, $\llbracket \Downarrow \text{uniform} \rrbracket_{CS} = \lambda = \llbracket \text{uniform} \rrbracket_{CS}$

Fix : Follows from the induction hypothesis and Lemma D.2

Abstraction : Follows directly from the induction hypothesis and Lemma D.3.

Application : Follows directly from the induction hypothesis and Lemma D.3.

Produce : Follows by inspection.

Force : Follows from the fact that the only well-typed closed programs of type $U\bar{\tau}$ are of the form $\text{thunk } t$ and from applying Lemma D.3 and the induction hypothesis.

Thunk : Follows by the induction hypothesis.

List Case : Using the distributivity of substitution, the fact that closed values of type list are either the empty list or the cons of a list, Lemma D.3 and the induction hypothesis, we can conclude.

Seq : This is the trickiest case. First, use the distributivity of substitution ($x \leftarrow t; u\{V_i/x_i\} = x \leftarrow t\{V_i/x_i\}; u\{V_i/x_i\}$). The proof follows from the induction hypothesis and the lemma above.

□

E EXAMPLES

E.1 Quickselect

Consider the Quickselect problem, which receives as input an unordered list l and a natural number n and outputs the the n -th largest element of l . This algorithm is very similar to quicksort: you choose a pivot, split the list into elements larger and smaller than it, then recurse on the appropriate

list.

```

 $\mu f : \text{list}(\mathbb{N}) \rightarrow \mathbb{N} \rightarrow F(\mathbb{N}).$ 
 $\lambda l : \text{list}(\mathbb{N}).$ 
 $\lambda n : \mathbb{N}.$ 
case  $l$  of
| nil  $\Rightarrow$ 
  produce nil
| (hd, tl)  $\Rightarrow$ 
   $len \leftarrow \text{length } l$ 
   $r \leftarrow \text{rand } len$ 
   $pivot \leftarrow len[r]$ 
   $(l_1, l_2) \leftarrow \text{biFilter } (\lambda n. \text{charge } 1; n \leq pivot) l$ 
   $lgth \leftarrow \text{len } l_1$ 
  if  $lgth < n - 1$  then
    (force  $f$ )  $l_1 n$ 
  elseif  $lgth == n - 1$  then
    produce  $pivot$ 
  else
    (force  $f$ )  $l_2 (n - lgth)$ 

```

In the best case, this algorithm runs in linear time and in the worst case, quadratic time. If we choose the pivot uniformly random, we get linear time, which is given by the following recurrence relation.

$$T(0) = 0$$

$$T(n) = n - 1 + \frac{1}{n} \sum_i T(i)$$

F PROOFS OF MISCELLANEOUS LEMMAS AND THEOREMS

F.1 Proof of Theorem 3.11

PROOF. Since $P_{\leq 1}$ is a monad, and the second component of the monad operations of $[0, \infty] \times P_{\leq 1}$ are identical to the ones of $P_{\leq 1}$, we only need to prove the monad laws for the first component. The unit laws follow from:

$$\pi_1(\eta^\#(r, \mu)) = r + 0 = r$$

$$\pi_1((f^\# \circ \eta)(x)) = \pi_1(f^\#(0, \delta_x)) = 0 + \pi_1(f(x))$$

While the last law requires a bit more work:

$$\pi_1((f^\# \circ g^\#)(r, \mu)) = \pi_1(f^\#(r + \int (\pi_1 \circ g) d\mu, (\pi_2 \circ g)^\#_{P_{\leq 1}}(\mu))) =$$

$$r + \int (\pi_1 \circ g) d\mu + \int (\pi_1 \circ f) d((\pi_2 \circ g)^\#_{P_{\leq 1}}(\mu)) = \pi_1((f^\# \circ g)^\#(r, \mu))$$

The last equation follows from the monad laws of $P_{\leq 1}$. □

F.2 Proof of Theorem 4.10

PROOF. The proof follows basically by commutativity of $P_{\leq 1}$:

$$\begin{aligned} \llbracket x \leftarrow t; y \leftarrow u; t' \rrbracket_{CS} &= \\ \int_{\mathbb{N} \times A} \int_{\mathbb{N} \times B} P_{\leq 1}(f)(\llbracket t' \rrbracket_{CS}(a, b)) \llbracket u \rrbracket_{CS}(dn_1 da) \llbracket t \rrbracket_{CS}(dn_2 db) &= \\ \int_{\mathbb{N} \times B} \int_{\mathbb{N} \times A} P_{\leq 1}(f)(\llbracket t' \rrbracket_{CS}(a, b)) \llbracket t \rrbracket_{CS}(dn_2 db) \llbracket u \rrbracket_{CS}(dn_1 da) &= \\ \llbracket y \leftarrow u; x \leftarrow t; t' \rrbracket_{CS}, \text{ where } f(n, c) = (n + n_1 + n_2, c) & \quad \square \end{aligned}$$

F.3 Proof of Theorem 5.9

PROOF. The proof follows by showing that there are monad structure morphisms $[0, \infty] \times (\mathbb{R} \rightarrow -) \rightarrow [0, \infty] \times T$ and $\varphi : [0, \infty] \times T \hookrightarrow K_{[0, \infty]}$. The proof is concluded by the uniqueness of factorizations.

The first monad morphism is $(id \times \psi)$ and the proof that this transformation is component-wise image-dense follows from the observation that the product order in $\omega\mathbf{Qbs}$ is given pairwise and ψ is, by assumption, dense in its image. The monad morphisms axioms follow from:

$$\begin{aligned} (id \times \psi)(\eta(a)) &= (id \times \psi)(0, \lambda r. a) = (0, \delta_a) = \eta(a) \\ (((id \times \psi) \circ f)^\# \circ (id \times \psi))(r, g) &= ((id \times \psi) \circ f)^\#(r, \psi(g), (\psi \circ \pi_2 \circ f)^\#(g)) = \\ (r + \int (\pi_1 \circ f) d(\psi(g)), (\psi \circ (\pi_2 \circ f)^\#)(g)) &= \\ (r + \int (\pi_1 \circ f \circ g) d(\lambda), (\psi \circ (\pi_2 \circ f)^\#)(g)) &= ((id \times \psi) \circ f^\#)(r, g) \end{aligned}$$

The second monad morphism has components $\varphi(r, \mu) = \lambda f. r + \int f d\mu$. The proof that this is indeed order reflecting is a direct consequence of $T \hookrightarrow K_{[0, \infty]}$ being order reflecting.

The proof that this is indeed a monad morphism follows from, once again, a series of direct calculations. □