

Characterising Renaming within OCaml's Module System: Theory and Implementation

Anonymous Author(s)

Abstract

We present an abstract, set-theoretic denotational semantics for a significant subset of OCaml and its module system in order to reason about the correctness of renaming value bindings. Our abstract semantics captures information about the binding structure of programs. Crucially for renaming, it also captures information about the relatedness of different declarations that is induced by the use of various different language constructs (e.g. functors, module types and module constraints). Correct renamings are precisely those that preserve this structure. We demonstrate that our semantics allows us to prove various high-level, intuitive properties of renamings. We also show that it is sound with respect to a (domain-theoretic) denotational model of the operational behaviour of programs. This formal framework has been implemented in a prototype refactoring tool for OCaml that performs renaming.

Keywords Adequacy, denotational semantics, dependencies, modules, module types, OCaml, refactoring, renaming, static semantics.

1 Introduction

Refactoring is the process of changing *how* a program works without changing *what* it does, and is a necessary and ongoing process in both the development and maintenance of any codebase [10]. Whilst individual refactoring steps are often conceptually very simple, applying them in practice can be complex, involving many repeated but subtly varying changes across the entire codebase. Moreover refactorings are, by and large, context sensitive, meaning that carrying them out by hand can be error-prone and the use of general-purpose utilities (even powerful ones such as `grep` and `sed`) is only effective up to a point.

This immediately poses a challenge, but also presents an opportunity. The challenge is how to ensure, or check, a proposed refactoring does not change the behaviour of the program (or does so only in very specific ways). The opportunity is that since refactoring is fundamentally a mechanistic process it is possible to automate it. Indeed, this is desirable in order to avoid human-introduced errors. Our aim in this paper is to outline how we might begin to provide a solution to the dual problem of specifying and verifying the correctness of refactorings and building correct-by-construction automated refactoring tools for OCaml [21, 30].

Renaming is a quintessential refactoring, and so it is on this that we focus as a first step. Specifically, we look at renaming the bindings of values in modules. One might very well be tempted to claim that, since we are in a functional setting, this is simply α -conversion (as in λ -calculus) and thus trivial. This is emphatically not the case. OCaml utilises language constructs, particularly in its module system, that behave in fundamentally different ways to traditional variable binders. Thus, to carry out renaming in OCaml correctly, one must take the meaning of these constructs into account. Some of the issues are illustrated by the following example.

```
module type Stringable = sig
  type t
  val to_string : t -> string
end
module Pair(X : Stringable)(Y : Stringable) = struct
  type t = X.t * Y.t
  let to_string (x, y) =
    (X.to_string x) ^ " " ^ (Y.to_string y)
end
module Int = struct
  type t = int
  let to_string i = int_to_string i
end
module String = struct
  type t = string
  let to_string s = s
end
module P = Pair(Int)(Pair(String)(Int)) ;;
print_endline (P.to_string (0, ("!", 1))) ;;
```

This program defines a functor `Pair` that takes two modules as arguments, which must conform to the `Stringable` module type. It also defines two structures `Int` and `String`. It then uses these as arguments in applications of `Pair`, the result of which is bound as the module `P`. Suppose that, for some reason, we wish to rename the `to_string` function in the module `Int`. To do so correctly, we must take the following into account.

(i) Since `Int` is used as the first argument to an application of `Pair`, the `to_string` member of `Pair`'s first parameter must be renamed.

(ii) The first parameter of `Pair` is declared to be of module type `Stringable`, so `to_string` in `Stringable` must be renamed; similarly for the second parameter, since `Int` is also used as the second argument in an application of `Pair`.

(iii) `String` is also used as an argument in an application of `Pair`, thus its `to_string` member must be renamed too.

(iv) An application of `Pair` is used as an argument to another such application, meaning that we also need to rename `to_string` in the body of `Pair` itself.

(v) Since **P** is bound to the result of applying **Pair**, we must then instances of **P.to_string**.

Thus, renaming the binding **Int.to_string** actually *depends* on renaming many other bindings in the program: failing to rename any one of them would result in the program being rejected by the compiler. Moreover, this is not simply an artifact of choosing to rename this particular binding; if we were to start with, say, **to_string** in **String** or **Stringable** we would still have to rename the same set of bindings. These bindings are all *mutually* dependent on each other. Consequently, the phenomenon we observe here is distinct from the notion of a refactoring pre-condition [32]. Note that although, in this example, it seemingly suffices to simply ‘find-and-replace’ all occurrences of **to_string**, this is not generally the case. If the example simply used **String** as the second argument to the (outer) application of **Pair**, then we would not have to rename the binding of **to_string** in the body of the functor.

The salient point in this example is that the various definitions and declarations that must be renamed are not simply references that resolve to a single instance of some syntactic construct in the program. On the contrary, they are themselves binding constructs, which can bind occurrences of identifiers elsewhere in the program. Nevertheless, as noted above, they are connected through certain syntactic constructions, albeit in a different sense to the notion of variable binding with which we are familiar from λ -calculus. Since here names matter, one way of viewing the situation might be to see the mutually dependent declarations (and their referents) all as instances of the same ‘free variable’ in the program. Free variables cannot be α -renamed, and so this view highlights the gap compared with an understanding of renaming based in the λ -calculus.

One objection to the foregoing analysis might be that the wide-reaching footprint of this refactoring indicates it is not really a renaming, or that it is, in some sense, ‘undesirable’. As to the former we would argue that, whilst the changes are extensive, the only syntactic operation that has occurred is to replace one identifier with another—surely, by definition, a renaming. Regarding the latter, other alternatives are indeed possible. One could, for example, localise the changes by introducing a new module expression in the applications of **Pair** that wraps the reference to the **Int** module and reintroduces a binding with the old name.

```

155 module P = Pair
156   (struct include Int let to_string = <<new_name>> end)
157   (Pair(String)
158     (struct include Int let to_string = <<new_name>> end))

```

The point here is not that we are trying to dictate *which* refactoring should be applied in any particular case, but that we are able to characterise precisely which changes of name are (not) refactorings. We can therefore provide a sound foundation for a refactoring tool enabling programmers to safely modify their code.

Our Contributions

In this paper, we propose a formal framework for reasoning about renaming in a significant subset of the OCaml language. We define an abstract semantics for programs in this subset, which captures particular aspects of the structure of programs relevant for renaming value bindings. This comprises *name-invariant* information about binding structure and dependencies between value binding constructs. We then define correctness of renamings in terms of the preservation of this structure. We show that our semantics constitutes a sensible abstraction by proving that it is sound with respect a denotational semantics of the operational behaviour of programs. We use our semantics to develop a *theory of renaming*, in which we characterise correct renamings in a natural and intuitive way and prove that they enjoy desirable (de)composition properties. Finally, we have built a prototype refactoring tool for the full OCaml language based on the concepts elucidated by our framework. We have evaluated our tool on two large real-world codebases.

We have formalised our framework and some of the renaming theory in the Coq proof assistant [38]. This is included as supplementary material with our submission. Results which have not yet been proved are marked as conjectures. We have also included as supplementary material an appendix containing a proof sketch of the adequacy result in section 5, and a high-level elaboration of proofs for the renaming theory.

While the paper describes the work in the context of OCaml modules, the approach can be used to understand aspects of (re)namings in other languages, such as Haskell (classes and instances), and Java (interfaces).

Paper Outline. In section 2, we describe the subset of OCaml that we study, and formally define operations that carry out renaming. We then present our abstract renaming semantics in section 3, before developing a formal theory of renaming in section 4. Section 5 shows that our renaming semantics is sound with respect to a denotational model of the operational behaviour of our calculus. In section 6 we describe our prototype refactoring tool and experimental evaluation. Section 7 surveys related work and section 8 concludes.

2 An OCaml Module Calculus

The subset of OCaml for which we build our formal theory is defined in fig. 1. It extends the calculus considered in [19, 20] and consists, essentially, of a two-level lambda calculus: the ‘core’ level defines basic values of the language (e.g. functions), whereas the other comprises the module system. The module system contains structures, functors, and module types (with module constraints and destructive module substitutions), along with **include** statements. Since value types do not interact with the renaming that we consider, we do not include a language for defining them. Thus, in order for our calculus to count as valid actual OCaml code, we use

Module Paths	Extended Module Paths	Value Expressions	Programs
$p ::= x \mid p.x$	$q ::= x \mid q.x \mid q(q)$	$e ::= v \mid p.v \mid \mathbf{let} \ v = e \ \mathbf{in} \ e \mid \mathbf{fun} \ v \rightarrow e \mid ee$	$P ::= e \mid \mathbf{module} \ x = m \ ; \ ; \ P$
Module Types $M ::= t \mid p.t \mid \mathbf{sig} \ S \ \mathbf{end} \mid \mathbf{functor} \ (x : M) \rightarrow M \mid M \ \mathbf{with} \ \mathbf{module} \ x = q \mid M \ \mathbf{with} \ \mathbf{module} \ x := q$			
Signature Body $S ::= \varepsilon \mid D \ ; \ ; \ S$	Signature Components $D ::= \mathbf{val} \ v : _ \mid \mathbf{module} \ x : M \mid \mathbf{module} \ \mathbf{type} \ t \mid \mathbf{module} \ \mathbf{type} \ t = M \mid \mathbf{include} \ M$		
Module Expressions $m ::= p \mid \mathbf{struct} \ s \ \mathbf{end} \mid \mathbf{functor} \ (x : M) \rightarrow m \mid m(m) \mid m : M$			
Structure Body $s ::= \varepsilon \mid d \ ; \ ; \ s$	Structure Components $d ::= \mathbf{let} \ v = e \mid \mathbf{module} \ x = m \mid \mathbf{module} \ \mathbf{type} \ t = M \mid \mathbf{include} \ m$		

Figure 1. Syntax of a core calculus for OCaml with modules.

OCaml's underscore syntax for anonymous type variables in value declarations in signatures, e.g. `sig val foo : _ end`.

Other features of OCaml's module system that we do not model, but which nonetheless interact with renaming, include: (local) open statements; recursive and first-class modules; module type extraction; and type-level module aliases. The first three should only require straightforward extensions of the approach we describe in this paper. Modelling type-level aliases correctly is more challenging, as they interact non-trivially with module type constraints [2].

We have assumed (disjoint) sets \mathcal{M} , \mathcal{T} , and \mathcal{V} of module, module type, and value identifiers, respectively. These are ranged over by x , t , and v , respectively, and we use l to range over the set $\mathcal{I} = \mathcal{M} + \mathcal{T} + \mathcal{V}$ of all identifiers. In real OCaml, both module identifiers and module type identifiers belong to the same lexical class. However, it will be convenient to distinguish them in our formalism. In any case it is syntactically unambiguous when such an identifier acts as a module identifier and when it acts as a module type identifier; thus we do not lose any generality in making this distinction.

2.1 Renaming Operations

To formalise the notion of carrying out renaming, we will take (fragments of) programs to be abstract syntax trees (ASTs). It will be convenient for us to consider ASTs as functions over some set \mathcal{L} of locations (ranged over by ℓ) returning local syntactic information. That is, for locations denoting internal nodes of the AST the function maps to the locations of the roots of the child subtrees and indicates which compound syntactic production is applied. For locations denoting leaves the function maps to the relevant identifier or constant. We will also assume that there is some *null* location $\perp \in \mathcal{L}$ that does not denote any location in any AST. This will be used by our semantics to indicate that a reference does not resolve to anything in a program. Although ASTs impose additional hierarchical structure on locations, we leave this implicit and do not further specify their concrete nature.

Definition 1. One program (fragment) σ' is the result of renaming another such σ , when: (i) $\text{dom}(\sigma) = \text{dom}(\sigma')$; (ii) $\sigma(\ell) \in \mathcal{V} \Leftrightarrow \sigma'(\ell) \in \mathcal{V}$; and (iii) if $\sigma(\ell) \notin \mathcal{V}$ then

$\sigma(\ell) = \sigma'(\ell)$. In this case, we call the pair (σ, σ') a *renaming* and write $\sigma \hookrightarrow \sigma'$.

That is, renaming is only allowed to replace value identifiers by other value identifiers, and must otherwise leave the program (fragment) unchanged.

We now define a number of syntactic concepts that will be useful in describing the action of renamings. Firstly, we consider the notion of the footprint of a renaming. This is all the locations in the program that are affected, or changed, by the renaming.

Definition 2 (Footprints). The *footprint* $\varphi(\sigma, \sigma')$ of a renaming $\sigma \hookrightarrow \sigma'$ is defined to be the set of locations (necessarily in both σ and σ') that are changed by the renaming: $\varphi(\sigma, \sigma') = \{\ell \mid \ell \in \text{dom}(\sigma) \wedge \sigma(\ell) \neq \sigma'(\ell)\}$. We write $\sigma \xrightarrow{\ell} \sigma'$ when ℓ is in the footprint of the renaming, and $\sigma \xrightarrow{v/\ell} \sigma'$ when moreover $\sigma'(\ell) = v$.

A general problem we are interested in is the following: given the location ℓ of some identifier in a program P and an identifier v that we wish to rename it to, can we produce a program P' such that $P \xrightarrow{v/\ell} P'$ is a *valid* renaming? Moreover, we are usually interested in finding such a P' that also *minimises* the footprint of the renaming. One purpose of the semantics that we define in section 3 is to enable us to provide solutions to this problem, as well as an effective abstraction of what constitutes validity for renaming.

Besides footprints, we are also interested in what we call the *dependencies* of a renaming. These are all the binding declarations modified by a renaming. In both the following definition and when presenting example syntax below, we will use subscripts on identifiers to indicate their unique position in the AST. In particular, numeric subscripts should not be taken to be part of the identifier itself.

Definition 3 (Declarations). The set $\text{decl}(\sigma)$ of (value) *declarations* in a program (fragment) σ is the set of all locations $\ell \in \text{dom}(\sigma)$ for which there exists $\ell' \in \text{dom}(\sigma)$ such that either: $\sigma(\ell') = \mathbf{val} \ v_\ell : _ ; ;$, $\sigma(\ell') = \mathbf{let} \ v_\ell = e ; ;$, $\sigma(\ell') = \mathbf{let} \ v_\ell = e \ \mathbf{in} \ e' ; ;$, or $\sigma(\ell') = \mathbf{fun} \ v_\ell \rightarrow e ; ;$.

Definition 4 (Dependencies). The *dependencies* $\delta(\sigma, \sigma')$ of $\sigma \hookrightarrow \sigma'$ are defined by $\delta(\sigma, \sigma') = \varphi(\sigma, \sigma') \cap \text{decl}(\sigma)$.

Intuitively, the dependencies should be the key piece of (syntactic) information required to characterise a renaming since we expect the remaining locations in the program that must be renamed to be simply those references that resolve to one of the dependencies.

We also formally define the references of a program (fragment) as follows.

Definition 5 (References). The set of (value) *references* of a program (fragment) σ is the set of locations $\ell \in \text{dom}(\sigma)$ such that $\sigma(\ell) \in \mathcal{V}$ and $\ell \notin \text{decl}(\sigma)$.

Notice that both the footprint and the dependencies of composite renamings are bounded by the footprints and dependencies, respectively, of their individual component renamings.

Proposition 1. For renamings $\sigma \hookrightarrow \sigma'$ and $\sigma' \hookrightarrow \sigma''$:

- (i) $\varphi(\sigma, \sigma'') \subseteq \varphi(\sigma, \sigma') \cup \varphi(\sigma', \sigma'')$.
- (ii) $\delta(\sigma, \sigma'') \subseteq \delta(\sigma, \sigma') \cup \delta(\sigma', \sigma'')$.

3 A Static Semantics for Renaming

In this section, we define a set-theoretic semantics for programs in our calculus that will allow us to reason about renaming values. The entities that comprise the meaning of a program are sets of (possibly nested) tuples of elements. Note that this allows us to also talk about functions, since these can be described by sets of ordered pairs. The semantics jointly describes binding resolution and dependency information in a name-invariant manner (using AST locations), and represents name-relevant information separately.

In the following presentation, we use standard notation for function update: i.e. $f[a \mapsto b]$ denotes the function that behaves like f except that $f(a) = b$. $f[a \mapsto b \mid a \in A]$ denotes the function that behaves like f except that $f(a) = b$ for all $a \in A$, and $f \setminus A$ the (partial) function that behaves like f but only has domain $\text{dom}(f) \setminus A$.

3.1 Semantic Elements

Our abstract semantics will consist of the following entities.

Binding Resolution is a function that maps the locations of uses of identifiers to binding instances of identifiers.

Definition 6 (Binding resolution). A binding resolution function \mapsto is a partial function between locations (we assume it does not map the null location \perp). We write $\ell \mapsto \ell'$ instead of $\mapsto(\ell) = \ell'$, and say that ℓ *resolves to* ℓ' .

The idea is that locations in the domain of the function will represent precisely the *references* in a program, and the function will describe the declaration that each reference resolves to.

Syntactic Characteristics that are captured by our semantics comprise the identifiers that are found at given locations.

This allows for the locations of binding instances of like identifiers to be related (cf. section 3.2 below).

Definition 7. A syntactic reification function $\rho : \mathcal{L} \rightarrow \mathcal{I}$ is a partial mapping from locations to identifiers (and we assume that ρ does not map the null location \perp). We write $\text{dom}_{\mathcal{V}}(\rho)$ to denote the set $\{\ell \mid \rho(\ell) \in \mathcal{V}\}$.

We can view syntactic reification functions as capturing a restricted view of ASTs, giving information only about those leaves that contain identifiers. The syntactic reification function can be used to give additional information, over and above the binding resolution function, about the declarations in a program (specifically, those which are never referenced).

Value Extensions capture sets of declarations that are all different facets of the same logical concept modelled in the program. For example, a program may contain many different functions named `compare` that act on values of various different data types, which might be related through the use of different signatures declaring values named `compare`, or the application of various functors to different modules. Although the different declarations may be distributed widely throughout the program, they all model a single concept or entity in the mind of the programmer or architecture of the system. These entities are high-level abstractions encoded via the global structure of program. When we rename a declaration, we must rename all parts of the program that constitute the logical entity of which it is part. The difficulty inherent in renaming OCaml arises since these high-level entities are not necessarily immediately evident, nor necessarily localised in the source code.

We call such collections of declarations the *extension*¹ of a high-level program abstraction. Ultimately, the extension is modelled by an equivalence class. However the structural relationships between the elements of an extension are more fine-grained and it is these that we capture, using a binary relation that we call a ‘kernel’. Taking the reflexive, symmetric and transitive closure of this kernel results in the equivalence relation whose equivalence classes we take to model extensions.

Definition 8. A value extension kernel \mathbb{E} is any binary relation on locations. $\hat{\mathbb{E}}$ denotes the reflexive, symmetric and transitive closure of \mathbb{E} .

For a location ℓ , we denote the $\hat{\mathbb{E}}$ -equivalence class containing ℓ by $[\ell]_{\hat{\mathbb{E}}}$. We also denote by $\mathcal{L}/_{\hat{\mathbb{E}}}$ the quotient of \mathcal{L} by $\hat{\mathbb{E}}$, i.e. the partitioning of the set of locations into $\hat{\mathbb{E}}$ -equivalence classes.

The notion of value extension will allow us to carry out renaming correctly by capturing the high-level, global structures present in a program. This is expressed in conjecture 2 below.

¹This is by analogy with Frege’s development and use of this term within the Logicist philosophical programme.

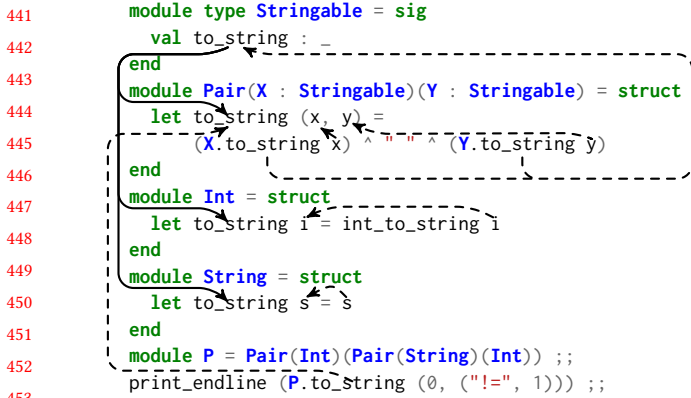


Figure 2. Graphical representation of the semantics.

To give an intuition as to how these elements are used, we show in fig. 2 a visual representation of the binding resolution function and value extension kernel that would be derived for the example in section 1. The binding resolution mappings are depicted using dashed arrows, and pairs in the value extension kernel by solid arrows.

3.2 Semantic Descriptions

In constructing the semantics of programs, we will need to keep track of the binding structure of modules and module types. We do so using *semantic descriptions*, which capture the locations of binding instances of identifiers and the nested structure of modules and module types. We distinguish two kinds of semantic descriptions: structural descriptions describe structures and signatures, while functorial descriptions describe functors and functor types.

Definition 9 (Semantic descriptions). Semantic descriptions, ranged over by Δ , are objects defined inductively as follows:

- A *component* is either: (i) a location ℓ ; or (ii) a pair of the form (ℓ, Δ) .
- A semantic description is either: (i) a set of components, in which case we call it *structural*; or (ii) a tuple of the form $((\ell, \Delta), \Delta')$, in which case we call it *functorial*.

We write \mathcal{D} for the set of all semantic descriptions. Additionally, we use D to range over structural descriptions only and will usually write functorial descriptions $((\ell, \Delta), \Delta')$ as $(\ell:\Delta)\rightarrow\Delta'$. We write $[D]$ to denote the set $\{\ell \mid \ell \in D\}$.

Basic components, comprising of simply a location, capture the locations of instances of identifiers bound to values. Components of the form (ℓ, Δ) represent subcomponents with further structure (i.e. sub-modules and sub-module types), along with the location of the instance of the identifier that they are bound to. Structural descriptions, which are sets of such components, thus describe the binding structure of structures and signatures. Functorial descriptions

$(\ell:\Delta)\rightarrow\Delta'$ capture that of functors and functor types: the left-hand member $\ell:\Delta$ captures the location of the parameter of the functor or functor type, along with a description of its declared type; the right-hand member of the pair, Δ' , describes the body.

We now define some operations on semantic descriptions that will be used to define the semantics of programs. In the following definitions, when we write $\rho(\ell)$ for a reification function ρ and a location ℓ , we mean this to also assert that ρ is defined on ℓ .

Superposition. We define a family of superposition operations \oplus_ρ on structural descriptions, parameterised by syntactic reification functions, that selectively combine the elements of the two descriptions based on syntactic information about locations contained in the reification function. The purpose of this is to sequentially combine the semantic description of two structure or signature fragments. In particular, it is used to model the effect of `include` statements.

Definition 10 (Description Superposition). The superposition operation \oplus_ρ on structural descriptions is defined by:

$$D \oplus_\rho D' = D' \cup \{\ell \mid \ell \in D \wedge \forall \ell' \in D'. \rho(\ell) \neq \rho(\ell')\} \\ \cup \{(\ell, \Delta) \mid (\ell, \Delta) \in D \wedge \forall (\ell', \Delta') \in D'. \rho(\ell) \neq \rho(\ell')\}$$

For example, consider the following modules.

```

module A = struct let foo1 = ...; let bar2 = ...; end
module B = struct include A let bar3 = ...; end

```

A semantic description of the module **A** consists of the set $D_A = \{1, 2\}$, while the remainder of the body of module **B** after the `include` statement consists of the set $D_{\text{body}} = \{3\}$. To form a description of the module **B**, we can superpose D_A and D_{body} with respect to the obvious reification function ρ that maps location 1 to `foo`, and locations 2 and 3 to `bar`. That is $D_B = D_A \oplus_\rho D_{\text{body}} = \{1, 3\}$. Here, the location 3 from D_{body} is chosen over 2 from D_A since ρ maps them both to the same identifier.

Join. We define a family of join operations \otimes_ρ on semantic descriptions, parameterised by syntactic reification functions, that each produce a value extension kernel from their input descriptions. The purpose of this operation is to extract the information about value extensions that is induced by the association of a module with a module type, either through an explicit module type annotation ($m : M$) or via a functor application ($m_1(m_2)$).

Definition 11 (Description Join). The description join operation \otimes_ρ is a binary operation on descriptions producing a value extension relation and is defined inductively as follows:

$$D_1 \otimes_\rho D_2 = \{(\ell_1, \ell_2) \mid \ell_1 \in D_1 \wedge \ell_2 \in D_2 \wedge \rho(\ell_1) = \rho(\ell_2)\} \\ \cup \{(\ell_1, \ell_2) \mid \exists (\ell, \Delta_1) \in D_1, (\ell', \Delta_2) \in D_2. \\ \rho(\ell) = \rho(\ell') \wedge (\ell_1, \ell_2) \in \Delta_1 \otimes_\rho \Delta_2\}$$

$$\begin{aligned}
& (\ell_1:\Delta_1)\rightarrow\Delta'_1 \otimes_\rho (\ell_2:\Delta_2)\rightarrow\Delta'_2 = (\Delta_1 \otimes_\rho \Delta_2) \cup (\Delta'_1 \otimes_\rho \Delta'_2) \\
& \Delta \otimes_\rho \Delta' = \emptyset \quad \text{otherwise}
\end{aligned}$$

To illustrate this, consider the functor from the example in section 1.

```

556 module type Stringable = sig val to_string1 : _ ;; end
557 module Pair = functor (X2 : Stringable) ->
558   functor (Y3 : Stringable) ->
559     struct let to_string4 = fun ...;; end

```

The **Stringable** module type is described by $D_{\text{Stringable}} = \{1\}$. For the **Pair** functor, our semantics constructs the description $D_{\text{Pair}} = (2:D_{\text{Stringable}})\rightarrow((3:D_{\text{Stringable}})\rightarrow\{5\})$. Applications of functors induce dependencies between the declarations in the type of the parameter, and corresponding bindings in the module used as the argument. These dependencies are computed by the join operation. Thus, for modules

```

567 module Int = struct let to_string5 = fun ...;; end
568 module String = struct let to_string6 = fun ...;; end

```

with descriptions $D_{\text{Int}} = \{5\}$ and $D_{\text{String}} = \{6\}$, respectively, an application **Pair**(**String**)(**Int**) induces dependencies $D_{\text{Stringable}} \otimes_\rho D_{\text{String}} = \{(1, 6)\}$ and $D_{\text{Stringable}} \otimes_\rho D_{\text{Int}} = \{(1, 5)\}$. Here, again, ρ is the obvious reification function (mapping 1 to `to_string`, 2 to `X`, etc.).

Modulation. We define a family of operations \blacktriangleright_ρ on semantic descriptions, also parameterised by a syntactic reification function, that produce another semantic description. This operation will be used to model how the description of a module is updated by a module type annotation.

Definition 12 (Description Modulation). The description modulation operation \blacktriangleright_ρ is a binary operation on semantic descriptions defined inductively as follows:

$$\begin{aligned}
D \blacktriangleright_\rho D' &= \{\ell \mid \ell \in D \wedge \exists \ell' \in D'. \rho(\ell) = \rho(\ell')\} \\
&\cup \{\ell' \mid \ell' \in D' \wedge \forall \ell \in D. \rho(\ell) \neq \rho(\ell')\} \\
&\cup \{(\ell, \Delta \blacktriangleright_\rho \Delta') \mid (\ell, \Delta) \in D \wedge \\
&\quad \exists \ell'. (\ell', \Delta') \in D' \wedge \rho(\ell) = \rho(\ell')\} \\
&\cup \{(\ell', \Delta') \mid (\ell', \Delta') \in D' \wedge \forall (\ell, \Delta) \in D. \rho(\ell) \neq \rho(\ell')\} \\
(\ell:\Delta_1)\rightarrow\Delta_2 \blacktriangleright_\rho (\ell':\Delta'_1)\rightarrow\Delta'_2 &= (\ell:(\Delta_1 \blacktriangleright_\rho \Delta'_1))\rightarrow(\Delta_2 \blacktriangleright_\rho \Delta'_2) \\
\Delta \blacktriangleright_\rho \Delta' &= \emptyset \quad \text{otherwise}
\end{aligned}$$

For example, consider the following module type, which is a weakening of the type of the **Pair** functor considered above.

```

597 module type Stringable2 = sig
598   val to_string7 : _ ;; val from_string8 : _ ;; end
599 module type WeakPair =
600   functor (X9 : Stringable2) ->
601     functor (Y10 : Stringable2) -> sig end

```

$D_{\text{Weak}} = (9:D_{\text{Stringable2}})\rightarrow((10:D_{\text{Stringable2}})\rightarrow\emptyset)$ describes the module type **WeakPair**, where $D_{\text{Stringable2}} = \{7, 8\}$. To describe the module $M = \text{Pair} : \text{WeakPair}$, we use the result

of the applying the modulation operation.

$$D_M = D_{\text{Pair}} \blacktriangleright_\rho D_{\text{Weak}} = (2:\{1, 8\})\rightarrow((3:\{1, 8\})\rightarrow\emptyset)$$

Notice that the result type has been restricted, but the types of the functor parameters in the original D_{Pair} description have been augmented by the additional `from_string` declarations (location 8) in the types of the parameters in D_{Weak} . Here, we intend that ρ has been updated with new mappings reflecting the identifiers occurring in **Stringable2** and **WeakPair** above.

We also define a family of selective modulation operations that modulate only certain elements of a structural description. This will be used to model the effect of a module constraint on a module type

Definition 13 (Selective Modulation). The selective modulation operation is a binary operation $\blacktriangleleft_\rho (x:\Delta')$ on semantic descriptions with respect to a module identifier, and is defined inductively as follows:

$$\begin{aligned}
D \blacktriangleleft_\rho (x:\Delta') &= \{\ell \mid \ell \in D\} \cup \{(\ell, \Delta) \mid (\ell, \Delta) \in D \wedge \rho(\ell) \neq x\} \\
&\quad \cup \{(\ell, \Delta \blacktriangleright_\rho \Delta') \mid (\ell, \Delta) \in D \wedge \rho(\ell) = x\} \\
(\ell:\Delta_1)\rightarrow\Delta_2 \blacktriangleleft_\rho (x:\Delta') &= \emptyset
\end{aligned}$$

For example, suppose we have the following module type.

```

630 module type Set = sig
631   module Elt11 : Stringable ;;
632   val empty12 : _ ;;
633 end

```

Consider also the following module.

```

635 Int2 = struct include Int; let from_string13 = ...; end

```

These can be described by $D_{\text{Set}} = \{12, (11, D_{\text{Stringable}})\}$ and $D_{\text{Int2}} = \{5, 13\}$. To compute the description of the module type given by **IntSet** = **Set with module Elt** = **Int2** we use selective modulation:

$$\begin{aligned}
D_{\text{IntSet}} &= D_{\text{Set}} \blacktriangleleft_\rho (\text{Elt} : D_{\text{Int2}}) \\
&= \{12, (11, (D_{\text{Stringable}} \blacktriangleright_\rho D_{\text{Int2}}))\} \\
&= \{12, (11, \{1, 13\})\}
\end{aligned}$$

Filtering. Lastly, we define an operation that removes elements from a structural description corresponding to a particular name, according to a given syntactic reification function. This will be used to model the effect of a destructive module substitution on a module type.

Definition 14 (Description Filtering). The function \setminus_ρ on semantic descriptions and (module) identifiers is defined by cases as follows:

$$\begin{aligned}
D \setminus_\rho x &= \{\ell \mid \ell \in D\} \cup \{(\ell, \Delta) \mid (\ell, \Delta) \in D \wedge \rho(\ell) \neq x\} \\
(\ell:\Delta)\rightarrow\Delta' \setminus_\rho x &= \emptyset
\end{aligned}$$

For example, to compute the description of the module type given by **IntSet2** = **Set with module Elt** := **Int2** we use filtering: $D_{\text{IntSet2}} = D_{\text{Set}} \setminus_\rho \text{Elt} = \{12\}$.

3.3 Semantic Environments

When constructing the semantics of programs, we will also need to keep track of the binding locations and descriptions of bound values, modules and module types. We do this using an *environment*, which is a pair (Γ_V, Γ_M) of functions $\Gamma_V : \mathcal{V} \rightarrow \mathcal{L}$ and $\Gamma_M : \mathcal{M} \cup \mathcal{T} \rightarrow \mathcal{D}$ that map value identifiers to the location in the program context to which they are bound, and map module and module type identifiers to semantic descriptions of the module or module type, respectively to which they are bound. We also require Γ_V to be injective on $\mathcal{L} \setminus \{\perp\}$, i.e. $\Gamma_V(v) = \Gamma_V(v') \neq \perp \Rightarrow v = v'$.

For notational convenience, we will write $\Gamma(v)$, $\Gamma(t)$, and $\Gamma(x)$ for $\Gamma_V(v)$, $\Gamma_M(t)$, and $\Gamma_M(x)$, respectively. Similarly, we will write $\Gamma[v \mapsto \ell]$, $\Gamma[t \mapsto \Delta]$, and $\Gamma[x \mapsto \Delta]$ for $(\Gamma_V[v \mapsto \ell], \Gamma_M)$, $(\Gamma_V, \Gamma_M[t \mapsto \Delta])$, and $(\Gamma_V, \Gamma_M[x \mapsto \Delta])$, respectively. Γ_\perp will denote the environment consisting of the functions that map every value identifier to the null location, and every module and module type identifier to the empty structural description (i.e. the empty set).

We say that a structural description D is *proper* for a reification function ρ when it satisfies: (i) $\rho(\ell) \in \mathcal{V}$ for all $\ell \in D$; (ii) $\rho(\ell) \in \mathcal{M} \cup \mathcal{T}$ for all $(\ell, \Delta) \in D$; and (iii) when $\ell, \ell' \in D$ or $(\ell, \Delta), (\ell', \Delta') \in D$ for distinct locations ℓ and ℓ' , then $\rho(\ell) \neq \rho(\ell')$. That is, each location in D corresponds to a *unique* identifier under ρ . In this case, we may treat it like a partial semantic environment and combine it with an existing environment Γ (written $\Gamma +_\rho D$) as follows:

$$(\Gamma +_\rho D)(\iota) = \begin{cases} \ell & \text{if } \ell \in D \text{ and } \rho(\ell) = \iota \\ \Delta & \text{if } (\ell, \Delta) \in D \text{ and } \rho(\ell) = \iota \\ \Gamma(\iota) & \text{otherwise} \end{cases}$$

3.4 Semantics of Programs

We will interpret programs as tuples $(\mapsto, \mathbb{E}, \rho)$ comprising a binding resolution function, a value extension kernel, and a syntactic reification function. We will use Σ to range over such tuples. We may also write Σ_{\mapsto} , $\Sigma_{\mathbb{E}}$, and Σ_ρ to denote the individual respective components of the semantics Σ .

To give the semantics of programs, we define two sorts of judgement, $\Sigma; \Gamma \vdash \sigma \rightsquigarrow \Sigma'$ and $\Sigma; \Gamma \vdash \sigma \rightsquigarrow (\Delta, \Sigma')$, which specify how the syntactic fragment σ extends the semantics Σ of a program context, described by Γ , to result in the semantics Σ' . The former sort of judgement applies when σ is a value expression, or a program (i.e. some number of module bindings followed by a value expression). The latter applies when σ is a module expression, module type expression, or the body of a structure or signature; in which case the judgement also derives a semantic description of σ .

The semantic judgements are defined inductively using the rules in fig. 3 below. They employ the following shorthand notation for specifying updates to $\Sigma = (\mapsto, \mathbb{E}, \rho)$.

- $\Sigma[\ell \mapsto \iota]$ stands for $(\mapsto, \mathbb{E}, \rho[\ell \mapsto \iota])$.
- $\Sigma[\ell \mapsto (\iota, \ell')]$ stands for $(\mapsto[\ell \mapsto \ell'], \mathbb{E}, \rho[\ell \mapsto \iota])$.

- $\Sigma[\Delta_1 \otimes \Delta_2]$ stands for $(\mapsto, \mathbb{E} \cup (\Delta_1 \otimes_\rho \Delta_2), \rho)$.
- $D \oplus_\Sigma D'$ stands for $D \oplus_\rho D'$.
- $\Delta \blacktriangleright_\Sigma \Delta'$ stands for $\Delta \blacktriangleright_\rho \Delta'$, and $\Delta \blacktriangleleft_\Sigma (x:\Delta')$ stands for $\Delta \blacktriangleleft_\rho (x:\Delta')$.
- $\Delta \setminus_\Sigma x$ stands for $\Delta \setminus_\rho x$.

Figure 3 elides the rules for standard module paths, since extended module paths are a strict superset of these. Moreover, for standard module paths, the judgement $\Sigma; \Gamma \vdash p \rightsquigarrow \Delta$ is used as a shorthand since, as can be straightforwardly determined, standard module paths do not update the semantics (although extended module paths, i.e. containing functor applications, do update the value extension kernel). We denote by Σ_\perp the *empty* semantics, i.e. the tuple consisting of the empty binding resolution function and syntactic reification functions and empty value extension kernel.

Under certain conditions (which we elide, but elaborate in the appendix), the semantics of fig. 3 are deterministic. Thus they allow us to interpret programs.

Definition 15 (Semantics of programs). We define families of (partial) interpretation functions $\llbracket \sigma \rrbracket_{\Sigma; \Gamma}$ and $\mathcal{D}_{\Sigma; \Gamma}(\sigma)$, indexed by pairs of semantics Σ and environments Γ , that return (when they exist) the unique Σ' and Δ , respectively, such that $\Sigma; \Gamma \vdash \sigma \rightsquigarrow \Sigma'$ or $\Sigma; \Gamma \vdash \sigma \rightsquigarrow (\Delta, \Sigma')$ holds. We write $\llbracket \sigma \rrbracket$ to mean $\llbracket \sigma \rrbracket_{\Sigma_\perp; \Gamma_\perp}$.

For a program P with $\llbracket P \rrbracket = \Sigma$, we will write \mapsto_P, \mathbb{E}_P , and ρ_P to mean $\Sigma_{\mapsto}, \Sigma_{\mathbb{E}}$, and Σ_ρ respectively.

The semantics naturally captures the syntactic information in a program pertaining to value identifiers.

Proposition 2. *If $\llbracket P \rrbracket$ is defined then $\text{ref}(P) = \text{dom}(\mapsto_P)$ and $\text{decl}(P) = \text{dom}_V(\rho_P) \setminus \text{dom}(\mapsto_P)$.*

An important point to note is that, even assuming an untyped value language, our semantics does *not* guarantee the well-typedness of programs. We consider this a feature rather than a bug since we see issues of renaming as orthogonal to type safety. Indeed, it is often desirable to be able to carry out renaming on incomplete (ill-typed) programs, and our semantics facilitates this. On the other hand, we can preserve well-typedness during renaming since the semantics captures the information required for renaming to also occur within module types. This also allows us to properly reason about renaming with respect to encapsulation, as illustrated by the following example.

```

module A = struct let foo = ...;; let bar = ...;; end
module B = struct
  include A : sig val foo : _ end;;
  let bar = ...;;
end

```

The **include** of module **A** in **B** is restricted by a module type, which serves to hide the fact that **A** contains a binding of **bar**. Thus, the binding of **bar** given in module **B** does not introduce any shadowing. The result is that we can rename **A.bar** and **B.bar** independently, whereas otherwise

(Extended) Module Paths

$$\frac{}{\Sigma; \Gamma \vdash x \rightsquigarrow (\Gamma(x), \Sigma)} \quad \frac{\Sigma; \Gamma \vdash q \rightsquigarrow (D, \Sigma')}{\Sigma; \Gamma \vdash q.x \rightsquigarrow (\Delta, \Sigma')} \quad (\exists \ell. \Sigma'_\rho(\ell) = x \wedge (\ell, \Delta) \in D) \quad \frac{\Sigma; \Gamma \vdash q_1 \rightsquigarrow ((\ell: \Delta_1) \rightarrow \Delta_2, \Sigma'') \quad \Sigma''; \Gamma \vdash q_2 \rightsquigarrow (\Delta'_1, \Sigma')}{\Sigma; \Gamma \vdash q_1 (q_2) \rightsquigarrow (\Delta_2, \Sigma'[\Delta_1 \otimes \Delta'_1])}$$

Value Expressions

$$\frac{}{\Sigma; \Gamma \vdash v_\ell \rightsquigarrow \Sigma[\ell \mapsto (v, \Gamma(v))]} \quad \frac{\Sigma; \Gamma \vdash p \rightsquigarrow D}{\Sigma; \Gamma \vdash p.v_\ell \rightsquigarrow \Sigma[\ell \mapsto (v, \ell')]} \quad (\Sigma_\rho(\ell') = v \wedge \ell' \in D) \quad \frac{\Sigma; \Gamma \vdash p \rightsquigarrow D}{\Sigma; \Gamma \vdash p.v_\ell \rightsquigarrow \Sigma[\ell \mapsto (v, \perp)]} \quad (\forall \ell' \in D. \Sigma_\rho(\ell') \neq v)$$

$$\frac{\Sigma; \Gamma \vdash e_1 \rightsquigarrow \Sigma'' \quad \Sigma''[\ell \mapsto v]; \Gamma[v \mapsto \ell] \vdash e_2 \rightsquigarrow \Sigma'}{\Sigma; \Gamma \vdash \mathbf{let} \ v_\ell = e_1 \ \mathbf{in} \ e_2 \rightsquigarrow \Sigma'} \quad \frac{\Sigma[\ell \mapsto v]; \Gamma[v \mapsto \ell] \vdash e \rightsquigarrow \Sigma'}{\Sigma; \Gamma \vdash \mathbf{fun} \ v_\ell \rightarrow e \rightsquigarrow \Sigma'} \quad \frac{\Sigma; \Gamma \vdash e_1 \rightsquigarrow \Sigma'' \quad \Sigma''; \Gamma \vdash e_2 \rightsquigarrow \Sigma'}{\Sigma; \Gamma \vdash e_1 e_2 \rightsquigarrow \Sigma'}$$

Signature Bodies

$$\frac{\Sigma[\ell \mapsto v]; \Gamma[v \mapsto \ell] \vdash S \rightsquigarrow (D, \Sigma')}{\Sigma; \Gamma \vdash \varepsilon \rightsquigarrow (\emptyset, \Sigma)} \quad \frac{}{\Sigma; \Gamma \vdash \mathbf{val} \ v_\ell : _ ; ; S \rightsquigarrow (\{\ell\} \oplus_{\Sigma'} D, \Sigma'[\{\ell\} \otimes D])}$$

$$\frac{\Sigma; \Gamma \vdash M \rightsquigarrow (\Delta, \Sigma'') \quad \Sigma''[\ell \mapsto x]; \Gamma[x \mapsto \Delta] \vdash S \rightsquigarrow (D, \Sigma')}{\Sigma; \Gamma \vdash \mathbf{module} \ x_\ell : M ; ; S \rightsquigarrow (\{\ell, \Delta\} \oplus_{\Sigma'} D, \Sigma')} \quad \frac{\Sigma; \Gamma \vdash M \rightsquigarrow (D, \Sigma'') \quad \Sigma''; \Gamma +_{\Sigma''} D \vdash S \rightsquigarrow (D', \Sigma')}{\Sigma; \Gamma \vdash \mathbf{include} \ M ; ; S \rightsquigarrow (D \oplus_{\Sigma'} D', \Sigma'[[D] \otimes D'])} \quad (D \text{ proper for } \Sigma''_\rho)$$

$$\frac{\Sigma[\ell \mapsto t]; \Gamma[t \mapsto \emptyset] \vdash S \rightsquigarrow (D, \Sigma')}{\Sigma; \Gamma \vdash \mathbf{module type} \ t_\ell ; ; S \rightsquigarrow (\{\ell, \emptyset\} \oplus_{\Sigma'} D, \Sigma')} \quad \frac{\Sigma; \Gamma \vdash M \rightsquigarrow (\Delta, \Sigma'') \quad \Sigma''[\ell \mapsto t]; \Gamma[t \mapsto \Delta] \vdash S \rightsquigarrow (D, \Sigma')}{\Sigma; \Gamma \vdash \mathbf{module type} \ t_\ell = M ; ; S \rightsquigarrow (\{\ell, \Delta\} \oplus_{\Sigma'} D, \Sigma')}$$

Module Types

$$\frac{}{\Sigma; \Gamma \vdash t \rightsquigarrow (\Gamma(t), \Sigma)} \quad \frac{\Sigma; \Gamma \vdash p \rightsquigarrow D}{\Sigma; \Gamma \vdash p.t \rightsquigarrow (\Delta, \Sigma)} \quad \left(\begin{array}{l} \exists \ell. \Sigma_\rho(\ell) = t \\ \wedge (\ell, \Delta) \in D \end{array} \right) \quad \frac{\Sigma; \Gamma \vdash S \rightsquigarrow (\Delta, \Sigma')}{\Sigma; \Gamma \vdash \mathbf{sig} \ S \ \mathbf{end} \rightsquigarrow (\Delta, \Sigma')}$$

$$\frac{\Sigma; \Gamma \vdash M_1 \rightsquigarrow (\Delta, \Sigma'') \quad \Sigma''[\ell \mapsto x]; \Gamma[x \mapsto \Delta] \vdash M_2 \rightsquigarrow (\Delta', \Sigma')}{\Sigma; \Gamma \vdash \mathbf{functor} \ (x_\ell : M_1) \rightarrow M_2 \rightsquigarrow ((\ell: \Delta) \rightarrow \Delta', \Sigma')}$$

$$\frac{\Sigma; \Gamma \vdash M \rightsquigarrow (\Delta, \Sigma'') \quad \Sigma''[\ell \mapsto x]; \Gamma \vdash q \rightsquigarrow (\Delta', \Sigma')}{\Sigma; \Gamma \vdash M \ \mathbf{with module} \ x_\ell = q \rightsquigarrow (\Delta \triangleleft_{\Sigma'} (x: \Delta'), \Sigma'[\Delta \otimes \{\ell, \Delta'\}])} \quad \frac{\Sigma; \Gamma \vdash M \rightsquigarrow (\Delta, \Sigma'') \quad \Sigma''[\ell \mapsto x]; \Gamma \vdash q \rightsquigarrow (\Delta', \Sigma')}{\Sigma; \Gamma \vdash M \ \mathbf{with module} \ x_\ell := q \rightsquigarrow (\Delta \setminus_{\Sigma'} x, \Sigma'[\Delta \otimes \{\ell, \Delta'\}])}$$

Structure Bodies

$$\frac{\Sigma; \Gamma \vdash e \rightsquigarrow \Sigma'' \quad \Sigma''[\ell \mapsto v]; \Gamma[v \mapsto \ell] \vdash s \rightsquigarrow (D, \Sigma')}{\Sigma; \Gamma \vdash \varepsilon \rightsquigarrow (\emptyset, \Sigma)} \quad \frac{\Sigma; \Gamma \vdash \mathbf{let} \ v_\ell = e ; ; s \rightsquigarrow (\{\ell\} \oplus_{\Sigma'} D, \Sigma'[\{\ell\} \otimes D])}{\Sigma; \Gamma \vdash \mathbf{let} \ v_\ell = e ; ; s \rightsquigarrow (\{\ell\} \oplus_{\Sigma'} D, \Sigma'[\{\ell\} \otimes D])} \quad \frac{\Sigma; \Gamma \vdash m \rightsquigarrow (D, \Sigma'') \quad \Sigma''; \Gamma +_{\Sigma''} D \vdash s \rightsquigarrow (D', \Sigma')}{\Sigma; \Gamma \vdash \mathbf{include} \ m ; ; s \rightsquigarrow (D \oplus_{\Sigma'} D', \Sigma'[[D] \otimes D'])} \quad (D \text{ proper for } \Sigma''_\rho)$$

$$\frac{\Sigma; \Gamma \vdash m \rightsquigarrow (\Delta, \Sigma'') \quad \Sigma''[\ell \mapsto x]; \Gamma[x \mapsto \Delta] \vdash s \rightsquigarrow (D, \Sigma')}{\Sigma; \Gamma \vdash \mathbf{module} \ x_\ell = m ; ; s \rightsquigarrow (\{\ell, \Delta\} \oplus_{\Sigma'} D, \Sigma')} \quad \frac{\Sigma; \Gamma \vdash M \rightsquigarrow (\Delta, \Sigma'') \quad \Sigma''[\ell \mapsto t]; \Gamma[t \mapsto \Delta] \vdash s \rightsquigarrow (D, \Sigma')}{\Sigma; \Gamma \vdash \mathbf{module type} \ t_\ell = M ; ; s \rightsquigarrow (\{\ell, \Delta\} \oplus_{\Sigma'} D, \Sigma')}$$

Module Expressions and Programs

$$\frac{\Sigma; \Gamma \vdash s \rightsquigarrow (\Delta, \Sigma')}{\Sigma; \Gamma \vdash \mathbf{struct} \ s \ \mathbf{end} \rightsquigarrow (\Delta, \Sigma')} \quad \frac{\Sigma; \Gamma \vdash m \rightsquigarrow (\Delta_1, \Sigma'') \quad \Sigma''; \Gamma \vdash M \rightsquigarrow (\Delta_2, \Sigma')}{\Sigma; \Gamma \vdash m : M \rightsquigarrow (\Delta_1 \blacktriangleright_{\Sigma'} \Delta_2, \Sigma'[\Delta_1 \otimes \Delta_2])} \quad \frac{\Sigma; \Gamma \vdash m_1 \rightsquigarrow ((\ell: \Delta_1) \rightarrow \Delta_2, \Sigma'') \quad \Sigma''; \Gamma \vdash m_2 \rightsquigarrow (\Delta'_1, \Sigma')}{\Sigma; \Gamma \vdash m_1 (m_2) \rightsquigarrow (\Delta_2, \Sigma'[\Delta_1 \otimes \Delta'_1])}$$

$$\frac{\Sigma; \Gamma \vdash M \rightsquigarrow (\Delta, \Sigma'') \quad \Sigma''[\ell \mapsto x]; \Gamma[x \mapsto \Delta] \vdash m \rightsquigarrow (\Delta', \Sigma')}{\Sigma; \Gamma \vdash \mathbf{functor} \ (x_\ell : M) \rightarrow m \rightsquigarrow ((\ell: \Delta) \rightarrow \Delta', \Sigma')} \quad \frac{\Sigma; \Gamma \vdash m \rightsquigarrow (\Delta, \Sigma'') \quad \Sigma''[\ell \mapsto x]; \Gamma[x \mapsto \Delta] \vdash P \rightsquigarrow \Sigma'}{\Sigma; \Gamma \vdash \mathbf{module} \ x_\ell = m ; ; P \rightsquigarrow \Sigma'}$$

Figure 3. The abstract renaming semantics of the OCaml calculus.

we would consider the latter to shadow the former and thus have to rename both together to preserve binding structure. A key feature of (module) types is that they should express such encapsulation properties.

4 Characterising Renaming

The primary purpose of our semantics is to distinguish ‘correct’ renamings from ‘incorrect’ ones. For example, given some declaration ℓ in program P and a new identifier v , it

might seem that $P' = P[\ell' \mapsto v \mid \ell' = \ell \vee \ell' \mapsto_P \ell]$ would be a good candidate for forming a minimal, valid renaming. That is, rename the identifier at location ℓ to v , as well as the identifiers at all the locations ℓ' that resolve to ℓ . As discussed in section 1 this is not always sufficient, and in general we find that we should modify multiple declarations and their associated references.

The first step, therefore, is to specify which renamings preserve meaning as captured by our semantics. The meaning

that we are interested in is *name-invariant* binding structure, which we capture at the semantic level via the following equivalence relations.

Definition 16 (Semantic Equivalence). We define the following equivalences on semantics and environments:

- $\Sigma \sim \Sigma'$ iff $\Sigma_{\rightarrow} = \Sigma'_{\rightarrow}$, $\Sigma_{\mathbb{E}} = \Sigma'_{\mathbb{E}}$, $\text{dom}(\Sigma_{\rho}) = \text{dom}(\Sigma'_{\rho})$, $\Sigma_{\rho}(\ell) \in \mathcal{V} \Leftrightarrow \Sigma'_{\rho}(\ell) \in \mathcal{V}$, and if $\Sigma_{\rho}(\ell) \notin \mathcal{V}$ then $\Sigma_{\rho}(\ell) = \Sigma'_{\rho}(\ell)$.
- $\Gamma \sim \Gamma'$ iff $\Gamma_{\mathcal{M}} = \Gamma'_{\mathcal{M}}$, and $\text{ran}(\Gamma_{\mathcal{V}}) = \text{ran}(\Gamma'_{\mathcal{V}})$.

When $\Sigma \sim \Sigma'$ and $\Gamma \sim \Gamma'$ hold, we write $(\Sigma, \Gamma) \sim (\Sigma', \Gamma')$.

Intuitively, this equivalence relation captures when two pairs of semantics and environments represent program contexts having the same binding structure regardless of the particular value identifiers that have been used. Notice that the equivalence relation on semantics comprises the same conditions on the syntactic reification function as are used to define renamings. With this equivalence we define what it means for a renaming to be valid.

Definition 17 (Valid Renamings). We say that a renaming $\sigma \hookrightarrow \sigma'$ is *valid with respect to* $\Sigma; \Gamma$, and write $\Sigma; \Gamma \vdash \sigma \hookrightarrow \sigma'$, when $\llbracket \sigma \rrbracket_{\Sigma; \Gamma}$ is defined, and there exists a semantics Σ' and environment Γ' with $(\Sigma', \Gamma') \sim (\Sigma, \Gamma)$ such that $\llbracket \sigma' \rrbracket_{\Sigma'; \Gamma'}$ is defined and $\llbracket \sigma \rrbracket_{\Sigma; \Gamma} \sim \llbracket \sigma' \rrbracket_{\Sigma'; \Gamma'}$. When $\Sigma_{\perp}; \Gamma_{\perp} \vdash \sigma \hookrightarrow \sigma'$ holds, then we simply say that the renaming $\sigma \hookrightarrow \sigma'$ is *valid*.

For whole programs, validity of renamings collapses to the following statement.

Proposition 3. $P \hookrightarrow P'$ is valid iff $\llbracket P \rrbracket$ and $\llbracket P' \rrbracket$ are defined and $\llbracket P \rrbracket \sim \llbracket P' \rrbracket$.

Thus, to check whether a renaming is valid, it suffices to compute the semantics of the original and renamed programs and then check that they are equivalent. We now proceed to explore some of the properties of valid renamings. That is to say, we begin to outline a theory of renaming for our OCaml calculus.

Firstly, as a basic sanity check, we note that renamings induce an equivalence relation on programs.

Proposition 4 (Equivalences). *The following properties hold:*

- i) $P \hookrightarrow P$ is a (valid) renaming (when $\llbracket P \rrbracket$ defined).
- ii) If $P \hookrightarrow P'$ is a (valid) renaming, then so is $P' \hookrightarrow P$.
- iii) If $P \hookrightarrow P'$ and $P' \hookrightarrow P''$ are (valid) renamings, then so is $P \hookrightarrow P''$.

A main objective for defining the semantics is to characterise renamings semantically. The following property shows that (up to unresolved references) a renaming is described by its dependencies and the binding resolution function.

Conjecture 1. *Suppose $P \hookrightarrow P'$ is a valid renaming, and let $L = \{\ell \mid \ell \in \delta(P, P') \vee \exists \ell' \in \delta(P, P'). \ell \mapsto_P \ell'\}$; then $L \subseteq \varphi(P, P')$ and $\ell \mapsto_P \perp$ for all $\ell \in \varphi(P, P') \setminus L$.*

This also means checking whether a renaming is invalid is cheaper than checking its validity, since we need only compute the semantics of the original program. Furthermore, the dependencies of a renaming are themselves characterised by the extension kernel.

Conjecture 2. *Let $P \hookrightarrow P'$ be a valid renaming, then $\delta(P, P')$ has a partitioning that is a subset of $\mathcal{L}_{/\mathbb{E}_P}$.*

The value extension kernel thus captures the dependencies inherent in a renaming: for a program P , all declarations belonging to an \mathbb{E}_P -equivalence class must be renamed together (along with their associated references), or none at all. In other words, *dependencies are value extensions*. This provides an alternative check for invalidity of renamings.

Given a declaration in a semantically meaningful program, it then follows from conjectures 1 and 2 that we can uniquely identify a lower bound for the footprint of any valid renaming containing the given declaration.

Conjecture 3. *For $P \xrightarrow{\ell} P'$ a valid renaming and $\ell \in \text{decl}(P)$, $\varphi(P, P') \supseteq \{\ell' \mid \ell' \in [\ell]_{\mathbb{E}_P} \vee \exists \ell'' \in [\ell]_{\mathbb{E}_P}. \ell' \mapsto_P \ell''\}$.*

This is, in fact, a tight bound since we can construct a valid renaming with exactly this footprint.

Proposition 5. *Suppose $\llbracket P \rrbracket$ is defined, $\ell \in \text{decl}(P)$, and $v \in \mathcal{V}$ does not occur in P , then $P \hookrightarrow P'$ is a valid renaming, where $P' = P[\ell' \mapsto v \mid \ell' \in [\ell]_{\mathbb{E}_P} \vee \exists \ell'' \in [\ell]_{\mathbb{E}_P}. \ell' \mapsto_P \ell'']$.*

Moreover, when a valid renaming does not have a minimal footprint, it is possible to decompose it into two, strictly smaller valid renamings.

Conjecture 4 (Factorisation). *Suppose $P \hookrightarrow P'$ is a valid renaming, and let ℓ and ℓ' be two distinct locations such that $\ell \in \varphi(P, P')$ and $\ell' \in \varphi(P, P')$, with $(\ell, \ell') \notin \mathbb{E}_P$; then there exists a P'' such that both $P \hookrightarrow P''$ and $P'' \hookrightarrow P'$ are valid, with $\varphi(P, P'') \subset \varphi(P, P')$ and $\varphi(P'', P') \subset \varphi(P, P')$.*

The reader may notice that our theory of renaming only utilises the equivalence relations induced by value extension kernels, rather than making any direct use of the structure of the value extension kernel itself. Nevertheless, we propose that our renaming theory could potentially make use of this detailed structure. One possibility is to define a complexity measure based on the ‘distance’ of the value extension kernel from its equivalence closure. We leave such investigations to future work.

5 Adequacy of the Semantics

The renaming semantics defined in section 3 leads to an intuitive theory for characterising renaming. However, it is also important that it constitutes a sensible abstraction of what we understand programs really to be. That is, the abstract semantics should be *adequate*, in the sense that it is a sound abstraction of the behavioural meaning of programs.

We now show that our renaming semantics is indeed adequate in this sense, by proving that if two renaming-related programs have equivalent abstract semantics then they have the same behaviour.

The model of program behaviour we consider is a denotational semantics that extends the model considered by Leroy in [20]. Our extensions cover the additional features of the module system incorporated by our OCaml calculus (i.e. `include` statements, module types as members of structures and signatures, and `module with` constraints on module types). However, we depart from that model in another important way: our model gives a denotational meaning to module types, which contribute towards the meaning of programs. This is because, as discussed in section 3 above, module types have meaning in the context of renaming. In contrast, the model of [20] simply ignores all module types in programs. For lack of space, we only describe the essential differences of our denotational model compared with [20]. The appendix contains the full definitions.

We assume an interpretation, using standard results, of value expressions (viz. lambda terms) in some domain \mathbb{F} containing an element `wrong` denoting run-time errors. We interpret modules in a domain \mathbb{M} satisfying:

$$\begin{aligned} \mathbb{M} &= \mathbb{D} + (\mathbb{M} \rightarrow \mathbb{M}) + \text{wrong} \\ \mathbb{D} &= (\mathcal{V} \rightarrow_{\text{fin}} \mathbb{F}) \times (\mathcal{T} \rightarrow_{\text{fin}} \mathbb{T}) \times (\mathcal{M} \rightarrow_{\text{fin}} \mathbb{M}) \end{aligned}$$

where \mathbb{T} is the set in which we interpret module types, defined inductively as the set X satisfying the following:

$$\begin{aligned} X &= D + (\mathcal{M} \times X) \times X + \text{wrong} \\ D &= \wp_{\text{fin}}(\mathcal{V}) \times (\mathcal{T} \rightarrow_{\text{fin}} X) \times (\mathcal{M} \rightarrow_{\text{fin}} X) \end{aligned}$$

The denotational semantics of programs is given by a function $\langle \cdot \rangle_{\theta}$, which interprets syntactic elements in their appropriate domains. As usual, it is parameterised by a denotational environment θ mapping identifiers to elements of the appropriate domain.

The interpretation of module types mirrors the way descriptions of module types are constructed by our abstract semantics. The main difference, then, between our denotational semantics and that of [20] is that module type denotations affect the meaning of modules. This happens in two ways. Firstly, the denotation of a module is modified by the denotation of a module type with which it is annotated.

$$\langle m : M \rangle_{\theta} = \text{let } d = \langle m \rangle_{\theta} \text{ in let } \tau = \langle M \rangle_{\theta} \text{ in } d : \tau$$

Here, we utilise a semantic operation $d : \tau$ on denotations d and τ , which essentially inserts ‘dynamic’ type checks. For example, if d denotes a structure containing some binding of v but τ denotes a signature not containing a declaration of v , then v will not be in the domain of $d : \tau$. In the reverse situation, v will be in the domain of $d : \tau$, but it will return `wrong` on being applied to v . This is analogous to the approach taken in gradual typing frameworks [36, 37], which insert casts that perform such dynamic checks.

Secondly, this operation is used to insert checks on the argument to a functor according to the module type declared for the corresponding parameter.

$$\begin{aligned} \langle \text{functor } (x : M) \rightarrow m \rangle_{\theta} &= \\ &\text{let } \tau = \langle M \rangle_{\theta} \text{ in } \lambda d. \langle m \rangle_{\theta[x \mapsto d : \tau]} \end{aligned}$$

We note that, for well-typed programs, this approach should be equivalent to the one ignoring all type annotations. Notwithstanding, by considering a ‘dynamically typed’ model we do not have to separately consider well-typedness.

Our abstract renaming semantics is sound with respect to the denotational semantics defined above. We write $\langle P \rangle$ to mean $\langle P \rangle_{\theta_{\perp}}$, where θ_{\perp} is the environment that maps everything to `wrong`.

Proposition 6 (Adequacy). $\langle P \rangle = \langle P' \rangle$ if $P \leftrightarrow P'$ is valid.

The converse result, completeness, does not hold. That is, there are renamings that preserve the operational meaning of programs, but which result in different abstract semantics. This is due to the fact that, according to our semantics, valid renamings must preserve all shadowing that occurs in programs. For example, consider the following contrived but nonetheless valid OCaml program.

```
module M = struct let foo = true let foo = 42 end
: sig val foo : bool val foo : int end ;;
M.foo ;;
```

Here there is shadowing in both the module expression and the module type. According to our semantics, the only valid renaming is the one that renames all instances of the identifier `foo`. However, it would be sufficient (in the sense that the result is denotationally equivalent) to rename both instances in the module type, but only the latter one in the module expression. It seems plausible that our semantics could be refined in order to reason about those cases in which (un)shadowing is allowed to occur, thus facilitating a completeness result. We leave this for future work.

6 ROTOR: A Refactoring Tool for OCaml

We have built a prototype refactoring tool for the OCaml language, called ROTOR (Reliable OCaml Tool for OCaml Refactoring), that carries out renaming based on the analysis modelled in our abstract semantics. The source code and a pre-compiled executable are available online [5, 6].

6.1 Implementation

The aim of our implementation was to produce a tool embodying proposition 5 above. That is, given a particular declaration in the input source code, the tool should produce a patch consisting of the minimal number of changes needed to correctly enact the renaming. In handling the OCaml language as a whole, we faced a number of challenges.

– In order to avoid having to build basic language processing functionality from scratch, we implemented ROTOR

in OCaml itself. This allowed us to reuse the compiler as a library, providing an abstract representation of the input source code directly. OCaml’s abstract syntax data type contains source code location information, which we used to produce accurate patches describing how to apply the renaming. We also relied on the recently developed visitors library [34] to automatically generate boilerplate code for traversing and processing the abstract syntax trees. This library provides similar functionality to that found in Haskell’s SYB [17] and Strafunski [18] libraries, or the Stratego/XT framework [9].

– For complex, real-world codebases the wider ecosystem and build pipeline of OCaml becomes relevant, as it introduces extra layers not present in the basic language itself. Two aspects of this were particularly relevant in implementing ROTOR. Firstly, OCaml has a preprocessor infrastructure called PPX [11]. This means that, in general, the abstract syntax that is processed by ROTOR may contain elements that do not correspond to actual source code. Moreover it is not always straightforward to determine when this is and is not the case, and our analysis must work on the post-processed code in order to fully compute the information it needs. Secondly, some build systems (e.g. dune [4]), in order to implement packaging and namespace separation, utilise custom mappings between the names of source files and the names of compiled modules, cf. [21, §8.12]. ROTOR must be aware of these custom mappings to be able to produce accurate patch information.

– The primary difficulty in implementing our analysis was computing the binding resolution and dependency information on which our analysis is built. Since it was not feasible to reimplement an entire binding analysis for the full language, we again relied on the OCaml compiler as much as possible. During type inference the compiler performs a binding analysis, assigning each binding a unique stamp. However, it only computes a partial view of the binding resolution function of our analysis. For value identifiers qualified by a module path (i.e. that refer to a binding inside another module), the compiler only provides the stamp of the outermost containing module whereas our binding resolution function provides the ‘stamp’ of the value binding itself.

For this reason, ROTOR approximates the abstract locations of our semantics using these logical paths. In fact, we had to extend the notion of paths implemented by the compiler, since they cannot refer to subcomponents of module types, or those of functors and their parameters. For each reference in the program, ROTOR can rely on information provided by the compiler to determine which logical path it resolves to. For each path, ROTOR must then compute the other paths it depends upon, i.e. which other declarations are in its value extension. It does this by comparing path prefixes whenever it encounters an `include` statement, module type annotation, module type constraint, or functor application. For example if, in analysing the dependencies of the

path `M.N.foo` (representing the `foo` value binding in the `N` submodule of module `M`), ROTOR encounters the module binding `module P = M : T`, it would generate dependencies on the paths `P.N.foo` and `T.N.foo`. An important point here is that, in our semantics, the logical paths `M.N.foo` and `P.N.foo` would denote the same (abstract) location, since module `P` is bound to module `M`. However, according to the information we can extract from the compiler, references might resolve to either of the paths. Thus, ROTOR must treat them as (logically) distinct dependencies.

ROTOR computes dependency information using a worklist algorithm, beginning with a working set containing just the path of the declaration to be renamed. For each dependency, it analyses the codebase to compute which other paths it depends upon, adding ones it has not previously processed to the working set. As each dependency is processed, ROTOR also identifies all of its references and builds up the final patch that can be applied to enact the renaming. At each point in the analysis, ROTOR checks to ensure that the new name does not introduce shadowing, or modify any shadowing that already occurs. If this is the case, ROTOR fails with a warning to the user. The renaming might also fail if ROTOR detects a declaration must be renamed that is not part of the input source code (e.g. a library function).

6.2 ROTOR in Practice

The aim of ROTOR is to provide a practical tool for refactoring “real world” OCaml code, but in doing this we have made a number of tradeoffs between the cost of handling certain features and the benefits that that would bring. We chose not to support modules that use PPX, because this can give rise to function declarations being automatically generated during PPX preprocessing; extending ROTOR to handle these cases would be very hard, as we would need to enable it to reason about meta-programming.

Other aspects – which lie outside core OCaml – include module type extraction; our choice here has been to concentrate on a set of language features that cover all essential aspects of the module system, such that other aspects could be treated using similar techniques.

We evaluated ROTOR on two substantial, real-world codebases. Firstly, Jane Street’s standard library overlay [15], comprising 869 source files in 77 libraries. Secondly, part of the OCaml (4.04.0) compiler itself [3] consisting of 502 source files. We analysed each codebase to extract its set of value bindings, which we used as test cases. For each case, we asked ROTOR to rename the binding to a fresh name not occurring in the codebase and tested the result by attempting to re-compile.

Setting aside the cases that we do not handle, and the cases which fail because they generate a requirement to rename an (external) library function, at the point of writing more than 70% of the tests pass; of the remainder, some are doubtlessly

1211 due to bugs, but others are due to the presence of features of
1212 the language so far unhandled by the system.

1213 As well as providing test data, this exercise has demon-
1214 strated the value of the dependency concept in practice.
1215 Among the refactorings for the OCaml compiler, more than
1216 thirty generate sets of dependencies of size at least 24, and
1217 over a hundred have non-trivial sets of dependencies. These
1218 more complex refactorings typically span multiple files, and
1219 generate multiple patches. In summary for the compiler, in
1220 the successful cases we have these data.

	Max	Mean	Mode
Files	19	3.8	3
Hunks	59	5.9	3
Dependencies	35	1.6	1
Avg. Hunks/File	15.0	1.5	1.0

1227 7 Related Work

1228 A general survey of refactoring research up until 2004 has
1229 been given by Mens and Tourwé [29]. Much work on refac-
1230 toring has been carried out within the object-oriented pro-
1231 gramming paradigm; a standard reference is [10]. Thompson
1232 and Li have carried out a survey of refactoring tools for func-
1233 tional languages [39] including the tools Wrangler [22, 23]
1234 (for Erlang [7]) and HaRe [24] (for Haskell [33]). Renaming,
1235 and perhaps refactoring generally, seems to be more difficult
1236 in a language like OCaml with its powerful module system.
1237 Erlang is dynamically typed, but has a flat module system,
1238 and Haskell, whilst possessing a powerful multi-feature type
1239 system, also does not support complex modules.

1240 It has long been recognised that, for correctness, refactor-
1241 ings generally require certain preconditions to hold [12]. As
1242 we have already noted, the notion of dependency that we
1243 describe in this paper is something other than a precondition
1244 and seems not to have been studied before. Our approach of
1245 constructing a semantic abstraction specifically for the pur-
1246 pose of refactoring, as far as we know, is also novel. It bears
1247 some similarity to work on program analysis via fact extrac-
1248 tion. This is the approach behind the codeQuest tool [13] and,
1249 more recently, the QL language [8] and Semmle platform
1250 [1]. The JunGL tool [41] uses this technique in the context
1251 of refactoring to check preconditions. However, these tools
1252 do not consider this technique as a semantic abstraction in
1253 a formal sense as we do. Lin and Holt consider an abstract
1254 formalization of fact extraction [26], and consider different
1255 notions of semantic completeness [27], but this is not tied to
1256 any language in particular and cannot obviously be applied
1257 to refactoring. Separately, Lin has also devised a (relational)
1258 algebraic procedure for binding resolution in various (imper-
1259 ative) languages, based on fact extraction [25]. Related to this
1260 is the recent work on scope graphs for name resolution [31]
1261 and static type checking [40]. This is a generic framework
1262 for specifying (and checking) static semantics of languages
1263 (including binding resolution), but does not present scope
1264

1265 graphs as abstractions of operational models. Menarini et
1266 al. take a semantic approach to code review, but do not ad-
1267 dress how semantics may guide automatic construction of
1268 refactorings [28].

1269 We have formally shown our renaming semantics to be an
1270 abstraction of an operational model of our OCaml calculus,
1271 which is an extension of the model considered in [19, 20] by
1272 Leroy. Rossberg et al. have also given a semantics for a large
1273 subset of OCaml and its module system via a translation
1274 to System F_ω [35]. However, since this translation requires
1275 programs to be well-typed, we did not follow this approach.
1276 The CakeML project [16] is a compiler stack for a large subset
1277 of OCaml that is formalised and fully verified in the HOL4
1278 theorem prover [14]. However, it currently contains only the
1279 most basic form of the module system.

1282 8 Conclusion

1283 In this paper we have presented a framework based on an
1284 abstract denotational semantics that allows us to reason
1285 about the correctness of renaming value bindings within
1286 OCaml modules. We have formally modelled a significant
1287 subset of the OCaml core language and its module system.
1288 Our abstract semantics allows us to characterise renamings
1289 which do not change the operational meaning of programs,
1290 and describe how they compose. A key concept that arose
1291 from our analysis was that of the *extension* of a value binding,
1292 this being the collection of bindings in the program that
1293 are related via the name-aware structures of the language.
1294 To the best of our knowledge, this is a novel concept not
1295 previously identified in the literature. We implemented our
1296 framework in a prototype tool called ROTOR, which is able
1297 to automatically carry out renaming on real-world OCaml
1298 code with a significant degree of success.

1299 **Future Work.** We would like to extend our approach to
1300 cover other features of OCaml's module system, such as
1301 first class and recursive modules, module type extraction,
1302 and type-level module aliases. We would also like to con-
1303 sider renaming module and module type bindings, as well
1304 as other kinds of refactoring. It will be interesting to see
1305 if our notion of value extension is flexible enough to cap-
1306 ture other language features and more complex refactorings.
1307 Our prototype tool, ROTOR, needs further development. It
1308 is our hope that it can become an industrially useful tool
1309 to the OCaml community. Furthermore, we would like to
1310 investigate whether our approach can be integrated into a
1311 mechanised formal framework, such as CakeML.

1314 References

- 1315 [1] [n. d.]. Semmle™. <https://www.semmle.com/> (accessed 11th November
1316 2018).
1317 [2] 2012. Mantis Bug Report 5514: “with module” semantics seem bro-
1318 ken. <https://caml.inria.fr/mantis/view.php?id=5514> Last accessed 9th
1319 Nov. 2018. Communicated to us by Leo White.

- [3] 2016. The Core OCaml System: Compilers, Runtime System, Base Libraries (version 4.04.0). <https://github.com/ocaml/ocaml/tree/4.04.0>
- [4] 2018. Dune: A Compsable Build System. <https://github.com/ocaml/dune>
- [5] 2018. A Prototype Refactoring Tool for OCaml. link removed for double blind review.
- [6] 2018. A Prototype Refactoring Tool for OCaml (Docker Image). link removed for double blind review.
- [7] Joe Armstrong, Robert Virding, Claes Wikström, and Mike Williams. 1996. *Concurrent Programming in ERLANG* (2nd ed.). Prentice Hall International (UK) Ltd., Hertfordshire, UK.
- [8] Pavel Avgustinov, Oege de Moor, Michael Peyton Jones, and Max Schäfer. 2016. QL: Object-oriented Queries on Relational Data. In *30th European Conference on Object-Oriented Programming, ECOOP 2016, July 18–22, 2016, Rome, Italy*. 2:1–2:25. <https://doi.org/10.4230/LIPIcs.ECOOP.2016.2>
- [9] Martin Bravenboer, Karl Trygve Kalleberg, Rob Vermaas, and Eelco Visser. 2008. Stratego/XT 0.17. A Language and Toolset for Program Transformation. *Sci. Comput. Program.* 72 (2008), 52–70. Issue 1–2. <https://doi.org/10.1016/j.scico.2007.11.003>
- [10] Martin Fowler, Kent Beck, John Brant, William Opdyke, and Don Roberts. 1999. *Refactoring: Improving the Design of Existing Code*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.
- [11] Alain Frisch. 2014. PPX and Extension Points. <https://lexifi.com/blog/ppx-and-extension-points> (blog post).
- [12] William G. Griswold and William F. Opdyke. 2015. The Birth of Refactoring: A Retrospective on the Nature of High-Impact Software Engineering Research. *IEEE Software* 32, 6 (2015), 30–38. <https://doi.org/10.1109/MS.2015.107>
- [13] Elnar Hajiyev, Mathieu Verbaere, Oege de Moor, and Kris De Volder. 2005. CodeQuest: Querying Source Code with Datalog. In *Companion to the 20th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2005, October 16–20, 2005, San Diego, CA, USA*. ACM, New York, NY, USA, 102–103. <https://doi.org/10.1145/1094855.1094884>
- [14] HOL4 2017. Interactive Theorem Prover. hol-theorem-prover.org
- [15] Jane Street. 2018. Standard Library Overlay. <https://github.com/janestreet/core>
- [16] Ramana Kumar, Magnus O. Myreen, Michael Norrish, and Scott Owens. 2014. CakeML: A Verified Implementation of ML. In *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14, San Diego, CA, USA, January 20–21, 2014*. ACM, New York, NY, USA, 179–192. <https://doi.org/10.1145/2535838.2535841>
- [17] Ralf Lämmel and Simon Peyton Jones. 2003. Scrap Your Boilerplate: A Practical Design Pattern for Generic Programming. In *Proceedings of TLDI'03: 2003 ACM SIGPLAN International Workshop on Types in Languages Design and Implementation, New Orleans, Louisiana, USA, January 18, 2003*. ACM, New York, NY, USA, 26–37. <https://doi.org/10.1145/604174.604179>
- [18] Ralf Lämmel and Joost Visser. 2003. A Strafunski Application Letter. In *Practical Aspects of Declarative Languages, 5th International Symposium, PADL 2003, New Orleans, LA, USA, January 13–14, 2003, Proceedings*. Springer-Verlag, Heidelberg Berlin, Germany, 357–375. https://doi.org/10.1007/3-540-36388-2_24
- [19] Xavier Leroy. 1994. Manifest Types, Modules, and Separate Compilation. In *Proceedings of POPL'94: 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Portland, Oregon, USA, January 17–21, 1994*. ACM, New York, NY, USA, 109–122. <https://doi.org/10.1145/174675.176926>
- [20] Xavier Leroy. 1995. Applicative Functors and Fully Transparent Higher-Order Modules. In *Conference Record of POPL'95: 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Francisco, California, USA, January 23–25, 1995*. ACM, New York, NY, USA, 142–153. <https://doi.org/10.1145/199448.199476>
- [21] Xavier Leroy, Damien Doligez, Alain Frisch, Jacques Garrigue, Didier Rémy, and Jérôme Vouillon. 2018. The OCaml System Release 4.07 Documentation and User's Manual. <http://caml.inria.fr/pub/docs/manual-ocaml/>
- [22] Huiqing Li and Simon J. Thompson. 2012. A Domain-Specific Language for Scripting Refactorings in Erlang. In *Fundamental Approaches to Software Engineering - 15th International Conference, FASE 2012, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2012, Tallinn, Estonia, March 24 – April 1, 2012, Proceedings*. 501–515. https://doi.org/10.1007/978-3-642-28872-2_34
- [23] Huiqing Li, Simon J. Thompson, George Orös, and Melinda Tóth. 2008. Refactoring with Wrangler, Updated: Data and Process Refactorings, and Integration with Eclipse. In *Proceedings of the 7th ACM SIGPLAN workshop on ERLANG, Victoria, BC, Canada, September 27, 2008*. 61–72. <https://doi.org/10.1145/1411273.1411283>
- [24] Huiqing Li, Simon J. Thompson, and Claus Reinke. 2005. The Haskell Refactorer, HaRe, and its API. *Electr. Notes Theor. Comput. Sci.* 141, 4 (2005), 29–34. <https://doi.org/10.1016/j.entcs.2005.02.053>
- [25] Yuan Lin. 2008. *Completeness of Fact Extractors and a New Approach to Extraction with Emphasis on the Refers-to Relation*. Ph.D. Dissertation. <http://hdl.handle.net/10012/3865>
- [26] Yuan Lin and Richard C. Holt. 2004. Formalizing Fact Extraction. *Electr. Notes Theor. Comput. Sci.* 94 (2004), 93–102. <https://doi.org/10.1016/j.entcs.2004.01.001>
- [27] Yuan Lin, Richard C. Holt, and Andrew J. Malton. 2003. Completeness of a Fact Extractor. In *10th Working Conference on Reverse Engineering, WCRE 2003, Victoria, Canada, November 13–16, 2003*. 196–205. <https://doi.org/10.1109/WCRE.2003.1287250>
- [28] Massimiliano Menarini, Yan Yan, and William G. Griswold. 2017. Semantics-assisted Code Review: An Efficient Toolchain and a User Study. In *Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering, ASE 2017, Urbana, IL, USA, October 30 – November 03, 2017*. IEEE Computer Society, 554–565. <https://doi.org/10.1109/ASE.2017.8115666>
- [29] Tom Mens and Tom Tourwé. 2004. A Survey of Software Refactoring. *IEEE Trans. Softw. Eng.* 30 (2004), 126–139. Issue 2. <https://doi.org/10.1109/TSE.2004.1265817>
- [30] Yaron Minsky, Anil Madhavapeddy, and Jason Hickey. 2013. *Real World OCaml: Functional Programming for the Masses*. O'Reilly Media, Sebastopol, CA, USA.
- [31] Pierre Neron, Andrew P. Tolmach, Eelco Visser, and Guido Wachsmuth. 2015. A Theory of Name Resolution. In *Programming Languages and Systems - 24th European Symposium on Programming, ESOP 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11–18, 2015, Proceedings*. 205–231. https://doi.org/10.1007/978-3-662-46669-8_9
- [32] William F. Opdyke. 1992. *Refactoring Object-Oriented Frameworks*. Ph.D. Dissertation. University of Illinois at Urbana-Champaign.
- [33] Simon Peyton Jones (Ed.). 2003. *Haskell 98 Language and Libraries: Revised Report*. Cambridge University Press, Cambridge, UK. haskell.org/onlinereport
- [34] François Pottier. 2017. Visitors Unchained. *PACMPL* 1, ICFP (2017), 28:1–28:28. <https://doi.org/10.1145/3110272>
- [35] Andreas Rossberg, Claudio V. Russo, and Derek Dreyer. 2014. F-ing Modules. *J. Funct. Program.* 24, 5 (2014), 529–607. <https://doi.org/10.1017/S0956796814000264>
- [36] Jeremy G. Siek and Walid Taha. 2006. Gradual Typing for Functional Languages. In *Scheme and Functional Programming 2006 - Proceedings of the 2006 Workshop on Scheme and Functional Programming, Portland, Oregon, Sunday September 17, 2006*, Robert Bruce Findler (Ed.). University of Chicago, 1100 East 58th Street, Chicago, IL 60637, 81–92. <https://newtraell.cs.uchicago.edu/research/publications/techreports/TR-2006-06> Technical Report TR-2006-06.

1431	[37] Jeremy G. Siek and Walid Taha. 2007. Gradual Typing for Objects. In <i>ECOOP 2007 - Object-Oriented Programming, 21st European Conference, Berlin, Germany, July 30 - August 3, 2007, Proceedings</i> . 2–27. https://doi.org/10.1007/978-3-540-73589-2_2	1486
1432		1487
1433		1488
1434	[38] The Coq Development Team. 2018. The Coq Proof Assistant, version 8.8.0. https://doi.org/10.5281/zenodo.1219885	1489
1435		1490
1436	[39] Simon Thompson and Huiqing Li. 2013. Refactoring Tools for Functional Languages. <i>Journal of Functional Programming</i> 23, 3 (2013), 293–350. https://doi.org/10.1017/S0956796813000117	1491
1437		1492
1438	[40] Hendrik van Antwerpen, Casper Bach Poulsen, Arjen Rouvoet, and Eelco Visser. 2018. Scopes As Types. <i>PACMPL</i> 2, OOPSLA (2018), 114:1–114:30. https://doi.org/10.1145/3276484	1493
1439		1494
1440		1495
1441	[41] Mathieu Verbaere, Ran Ettinger, and Oege de Moor. 2006. JunGL: A Scripting Language for Refactoring. In <i>28th International Conference on Software Engineering (ICSE 2006), Shanghai, China, May 20-28, 2006</i> . ACM, New York, NY, USA, 172–181. https://doi.org/10.1145/1134311	1496
1442		1497
1443		1498
1444		1499
1445		1500
1446		1501
1447		1502
1448		1503
1449		1504
1450		1505
1451		1506
1452		1507
1453		1508
1454		1509
1455		1510
1456		1511
1457		1512
1458		1513
1459		1514
1460		1515
1461		1516
1462		1517
1463		1518
1464		1519
1465		1520
1466		1521
1467		1522
1468		1523
1469		1524
1470		1525
1471		1526
1472		1527
1473		1528
1474		1529
1475		1530
1476		1531
1477		1532
1478		1533
1479		1534
1480		1535
1481		1536
1482		1537
1483		1538
1484		1539
1485		1540

A Proofs

Here we elaborate on the results stated in the main body of the paper, and provide proofs of those results that are not included in the Coq formalisation.

A.1 The Abstract Renaming Semantics

It was stated in section 3 that, under certain conditions, the semantics are deterministic. Here, we give the formal statement of this property.

We first have to define a notion of well-behavedness for semantic descriptions and environments. Given an interpretation of locations as identifiers (i.e. a syntactic reification function), a semantic description is well-behaved when each location in a (possibly nested) structural description corresponds to an identifier that is *unique* within that description.

Definition 18 (Well-behaved Descriptions). We define the subset of semantic descriptions that are *well-behaved* with respect to a given syntactic reification function ρ as the smallest set satisfying the following.

- A structural description D is well-behaved w.r.t. ρ when:
 - (i) $\ell \in D$ implies $\ell \in \text{dom}(\rho)$ and $\rho(\ell) \in \mathcal{V}$;
 - (ii) $(\ell, \Delta) \in D$ implies $\ell \in \text{dom}(\rho)$, $\rho(\ell) \in \mathcal{M} \cup \mathcal{T}$ and Δ is well-behaved w.r.t. ρ ; and
 - (iii) if $\rho(\ell) = \rho(\ell')$ for $\ell, \ell' \in D$ or $(\ell, \Delta), (\ell', \Delta') \in D$, then also $\ell = \ell'$.
- A functorial description $(\ell:\Delta) \rightarrow \Delta'$ is well-behaved w.r.t. ρ when both Δ and Δ' are well-behaved w.r.t. ρ .

That is, a semantic description that is well-behaved for ρ is proper for ρ ‘all the way down’.

We say that an environment Γ is well-behaved for a syntactic reification function ρ when $\Gamma(v) = \ell$ implies $\rho(\ell) = v$ for every $\ell \neq \perp$, and each Δ_i such that $\Gamma(i) = \Delta_i$ ($i \in \mathcal{M} \cup \mathcal{T}$) is well-behaved w.r.t. ρ . We say that an environment Γ or semantic description Δ is well-behaved for a *semantics* Σ when it is well-behaved w.r.t. the reification function ρ for which $\rho(i) = \ell$ if and only if $\Sigma_\rho(i) = \ell$ and $\ell \notin \text{dom}(\Sigma_{\rightarrow})$. We denoted by $\text{ran}_I(\Gamma)$ the set $\text{ran}(\Gamma_V) \cup \{\ell \mid \exists \Delta. (\ell, \Delta) \in \text{ran}(\Gamma_M)\}$.

Lemma 5 (Determinism). *For any program fragment σ , semantics Σ , and environment Γ that is well-behaved for Σ and satisfies $(\text{dom}(\Sigma_\rho) \cup \text{ran}_I(\Gamma)) \cap \text{dom}(\sigma) = \emptyset$, there is at most one Σ' and one Δ such that $\Sigma; \Gamma \vdash \sigma \rightsquigarrow \Sigma'$ or $\Sigma; \Gamma \vdash \sigma \rightsquigarrow (\Delta, \Sigma')$.*

Proof. Given in the Coq formalisation. By induction on the definition of the semantics. In fact, we need to use a stronger hypothesis involving the following additional invariants:

- (1) Σ' contains only locations in $\text{dom}(\Sigma_\rho)$ and $\text{dom}(\sigma)$;
- (2) Γ is well-behaved also for Σ' ;
- (3) for judgements $\Sigma; \Gamma \vdash \sigma : (\Delta, \Sigma')$, then Δ is well-behaved for Σ' ; and
- (4) Δ well-behaved w.r.t. Σ implies Δ well-behaved w.r.t. Σ' , for all Δ . \square

Thus, we specify that $\llbracket \sigma \rrbracket_{\Sigma; \Gamma}$ and $\mathcal{D}_{\Sigma; \Gamma}(\sigma)$ are only defined when Γ is well-behaved for Σ and $(\text{dom}(\Sigma_\rho) \cup \text{ran}_I(\Gamma)) \cap \text{dom}(\sigma) = \emptyset$. A consequence of lemma 5 is that (when defined) $\mathcal{D}_{\Sigma; \Gamma}(\sigma)$ is well-behaved w.r.t. ρ , where $\llbracket \sigma \rrbracket_{\Sigma; \Gamma} = (\rightarrow, \mathbb{E}, \rho)$.

The following property is necessary for a semantics to correspond to an actual program fragment.

Definition 19 (Properness). A semantics $\Sigma = (\rightarrow, \mathbb{E}, \rho)$ is called *proper* when it satisfies the following conditions.

- (i) $\text{dom}(\rightarrow) \cap \text{ran}(\rightarrow) = \emptyset$.
- (ii) $\ell \rightarrow \ell'$ and $\ell' \neq \perp$ implies $\rho(\ell) = \rho(\ell')$.
- (iii) $\rho(\ell) \in \mathcal{V}$, for all $\ell \in \text{dom}(\rightarrow) \cup \text{ran}(\rightarrow)$ with $\ell \neq \perp$.
- (iv) $\rho(\ell) = \rho(\ell') \in \mathcal{V}$, $\ell \notin \text{dom}(\rightarrow)$ and $\ell' \notin \text{dom}(\rightarrow)$, for all (ℓ, ℓ') in \mathbb{E} .

Note that the empty semantics is trivially proper. We can show that properness is preserved by the semantics.

Lemma 6. *Let Σ be proper, and environment Γ be well-behaved for Σ ; if $\Sigma; \Gamma \vdash \sigma \rightsquigarrow \Sigma'$ or $\Sigma; \Gamma \vdash \sigma \rightsquigarrow (\Delta, \Sigma')$ holds then Σ' is proper.*

Proof. By induction on the semantic rules. Given in the Coq formalisation. \square

The semantic characterisation of the syntactically defined references and declarations given in proposition 2 is a special case of the following lemma. We write $\text{decl}(\Sigma)$ to denote the set $\text{dom}_V(\Sigma_\rho) \setminus \text{dom}(\Sigma_{\rightarrow})$.

Proposition 7. *If $\llbracket \sigma \rrbracket_{\Sigma; \Gamma} = \Sigma'$ then:*

- (i) $\text{ref}(\sigma) = \text{dom}(\Sigma'_{\rightarrow}) \setminus \text{dom}(\Sigma_{\rightarrow})$.
- (ii) $\text{decl}(\sigma) = \text{decl}(\Sigma') \setminus \text{decl}(\Sigma)$.

Proof. By induction on the semantic rules. Given in the Coq formalisation. \square

We now justify the statement of validity for whole program renamings.

Proposition 3. *$P \hookrightarrow P'$ is valid iff $\llbracket P \rrbracket$ and $\llbracket P' \rrbracket$ are defined and $\llbracket P \rrbracket \sim \llbracket P' \rrbracket$.*

Proof. Notice that trivially Γ_\perp is well-behaved for Σ_\perp and, when restricting to pairs (Σ, Γ) such that Γ is well-behaved for Σ , we have $[(\Sigma_\perp, \Gamma_\perp)]_- = \{(\Sigma_\perp, \Gamma_\perp)\}$, whence the statement follows directly from definition 17. \square

We now consider some properties pertaining to the structure of the semantics and descriptions synthesised by the semantic rules. In an abuse of notation, we will write $\mathcal{L}(\Delta)$ to denote the set of all locations appearing in (a subcomponent) of Δ . For an environment Γ and identifier $i \in \mathcal{M} \cup \mathcal{T}$, we then write $\Gamma_{\mathcal{D}}(i)$ for the description Δ such that there exists ℓ with $\Gamma(i) = (\ell, \Delta)$, and $\text{ran}_{\mathcal{D}}(\Gamma)$ for the set $\bigcup_{i \in \mathcal{M} \cup \mathcal{T}} \mathcal{L}(\Gamma_{\mathcal{D}}(i))$.

Lemma 7. *If $\Sigma; \Gamma \vdash \sigma \rightsquigarrow (\Delta, \Sigma')$ then $\mathcal{L}(\Delta) \subseteq \text{dom}(\sigma) \cup \text{ran}_{\mathcal{D}}(\Gamma)$.*

Proof. By induction on the semantic rules. Included in the Coq formalisation. \square

Lemma 8. *If $\Sigma; \Gamma \vdash \sigma \rightsquigarrow (\Delta, \Sigma')$ then $\mathbb{E}' \setminus \mathbb{E} \subseteq L \times L$, for $L = \text{dom}(\sigma) \cup \text{ran}_{\mathcal{D}}(\Gamma)$, where \mathbb{E} and \mathbb{E}' are the extension kernels of Σ and Σ' , respectively.*

Proof. By induction on the semantic rules. Included in the Coq formalisation. \square

The following concepts of inclusion, relevance and freshness for semantics are central to proving many of the results in this paper. To express these properties, we use the following notation for partial functions f and g , and (binary) relation R :

- $f \subseteq g$ denotes that, for all $x \in \text{dom}(f)$, if $f(x) = y$ then $g(x) = y$;
- $f \setminus g$ denotes the function defined by $(f \setminus g)(x) = y$ if and only if $f(x) = y$ and either $g(x)$ undefined or $g(x) \neq y$;
- $f \subseteq R$ denotes that $f(x) = y$ only if $(x, y) \in R$.

The use of set-theoretic notation here is justified by the view of (partial) functions as sets of mappings (i.e. pairs). Notice that the following property holds.

Proposition 8. *Suppose that $f(x) = y$ but $x \notin \text{dom}(f \setminus g)$, then $g(x) = y$.*

Proof. Suppose, for contradiction, that in fact $g(x)$ is undefined or else $g(x) \neq y$. But then from the assumption that $f(x) = y$ we have, by definition, that $(f \setminus g)(x) = y$, which contradicts the assumption that $x \notin \text{dom}(f \setminus g)$. \square

The definitions of inclusion, relevance and freshness are as follows.

Definition 20 (Inclusion). We say that Σ' *includes* Σ , and write $\Sigma \subseteq \Sigma'$, when the following hold: (1) $\Sigma_{\rightsquigarrow} \subseteq \Sigma'_{\rightsquigarrow}$; (2) $\Sigma_{\mathbb{E}} \subseteq \Sigma'_{\mathbb{E}}$; and (3) $\Sigma_{\rho} \subseteq \Sigma'_{\rho}$. When we additionally have $\ell \in \text{dom}(\Sigma_{\rho}) \setminus \text{dom}(\Sigma_{\rightsquigarrow})$ implies $\ell \in \text{dom}(\Sigma'_{\rho}) \setminus \text{dom}(\Sigma'_{\rightsquigarrow})$ for all locations $\ell \in \mathcal{L}$, we say that Σ' *properly includes* Σ .

Definition 21 (Relevance). For semantics Σ and Σ' , and a set of locations $L \subseteq \mathcal{L}$, we say Σ' is *relevant for L over Σ* , and write $\Sigma' \setminus \Sigma \subseteq L$, when the following hold:

- (1) $\Sigma'_{\rightsquigarrow} \setminus \Sigma_{\rightsquigarrow} \subseteq L \times (L \cup \text{dom}(\Sigma_{\rho}))$
- (2) $\Sigma'_{\rho} \setminus \Sigma_{\rho} \subseteq L \times \mathcal{I}$
- (3) $\Sigma'_{\mathbb{E}} \setminus \Sigma_{\mathbb{E}} \subseteq (L \cup \text{dom}(\Sigma_{\rho}))^2 \setminus \text{dom}(\Sigma_{\rho})^2$

Definition 22 (Freshness). We say that a set $L \subseteq \mathcal{L}$ of locations is *fresh* for a semantics $\Sigma = (\rightsquigarrow, \mathbb{E}, \rho)$ when the following properties hold for all locations $\ell \in \mathcal{L}$:

- (1) $\ell \in \text{dom}(\rightsquigarrow) \cup \text{ran}(\rightsquigarrow) \Rightarrow \ell \notin L$
- (2) $(\exists \ell'. (\ell, \ell') \in \mathbb{E} \vee (\ell', \ell) \in \mathbb{E}) \Rightarrow \ell \notin L$
- (3) $\ell \in \text{dom}(\rho) \Rightarrow \ell \notin L$

For proper semantics, these properties are guaranteed by the interpretation function.

Lemma 9. *If $\llbracket \sigma \rrbracket_{\Sigma; \Gamma} = \Sigma'$ with Σ proper then: (1) Σ' properly includes Σ ; (2) Σ' is relevant for $\text{dom}(\sigma)$ over Σ ; and (3) $\text{dom}(\sigma)$ is fresh for Σ .*

Proof. Given in the Coq formalisation. The freshness property follows from properness and the preconditions for $\llbracket \sigma \rrbracket_{\Sigma; \Gamma}$ to be defined (cf. definition 15)—namely that $\text{dom}(\Sigma_{\rho}) \cap \text{dom}(\sigma) = \emptyset$. The other properties are shown by induction on syntactic structure. \square

Thus, the major utility of definitions 20 to 22 lies in the following result.

Lemma 10. *Take semantics $\Sigma_1, \Sigma_2, \Sigma'_1$ and Σ'_2 , with a set of locations $L \subseteq \mathcal{L}$ such that the following conditions hold:*

- $\Sigma_1 \subseteq \Sigma'_1$, and $\Sigma_2 \subseteq \Sigma'_2$;
- $\Sigma'_1 \setminus \Sigma_1 \subseteq L$ and $\Sigma'_2 \setminus \Sigma_2 \subseteq L$; and
- L is fresh for both Σ_1 and Σ_2 .

Then $\Sigma'_1 \sim \Sigma'_2$ implies that $\Sigma_1 \sim \Sigma_2$.

Proof. Let $\Sigma_1 = (\rightsquigarrow_1, \mathbb{E}_1, \rho_1)$, $\Sigma_2 = (\rightsquigarrow_2, \mathbb{E}_2, \rho_2)$, with $\Sigma'_1 = (\rightsquigarrow'_1, \mathbb{E}'_1, \rho'_1)$, and $\Sigma'_2 = (\rightsquigarrow'_2, \mathbb{E}'_2, \rho'_2)$. Since $\Sigma'_1 \sim \Sigma'_2$, we have by definition 16 that $\rightsquigarrow'_1 = \rightsquigarrow'_2$, $\mathbb{E}'_1 = \mathbb{E}'_2$, $\text{dom}(\rho'_1) = \text{dom}(\rho'_2)$, $\rho'_1(\ell) \in \mathcal{V} \Leftrightarrow \rho'_2(\ell) \in \mathcal{V}$, and $\rho'_1(\ell) = \rho'_2(\ell)$ if $\rho'_1(\ell) \notin \mathcal{V}$. We must show the following:

$(\rightsquigarrow_1 = \rightsquigarrow_2)$: To see that $\rightsquigarrow_1 \subseteq \rightsquigarrow_2$, take $(\ell, \ell') \in \rightsquigarrow_1$. Since $\Sigma_1 \subseteq \Sigma'_1$ it follows that $\rightsquigarrow_1 \subseteq \rightsquigarrow'_1$, and thus that $(\ell, \ell') \in \rightsquigarrow'_1$. Moreover, since $\rightsquigarrow'_1 = \rightsquigarrow'_2$ it then follows that $(\ell, \ell') \in \rightsquigarrow'_2$. Now, since L is fresh for Σ_1 , we have that $\ell \notin L$ and therefore, since Σ'_2 is relevant for L over Σ_2 , it follows that $\ell \notin \text{dom}(\rightsquigarrow'_2 \setminus \rightsquigarrow_2)$. However, since we have that $(\ell, \ell') \in \rightsquigarrow'_2$, by proposition 8 it must be that $(\ell, \ell') \in \rightsquigarrow_2$ as required. A symmetric chain of reasoning shows that $\rightsquigarrow_2 \subseteq \rightsquigarrow_1$, hence we conclude.

$(\mathbb{E}_1 = \mathbb{E}_2)$: To see that $\mathbb{E}_1 \subseteq \mathbb{E}_2$, take $(\ell, \ell') \in \mathbb{E}_1$ and reason as above that $(\ell, \ell') \in \mathbb{E}'_2$. Since, L is fresh for Σ_1 , it follows that neither $\ell \in L$ nor $\ell' \in L$. Then, since Σ'_2 is relevant for L over Σ_2 , we have by clause (3) of definition 21 that for any $(\ell_1, \ell_2) \in \mathbb{E}'_2 \setminus \mathbb{E}_2$ it must be that either $\ell_1 \in L$ or $\ell_2 \in L$. Thus, $(\ell, \ell') \notin \mathbb{E}'_2 \setminus \mathbb{E}_2$. Therefore, since $(\ell, \ell') \in \mathbb{E}'_2$, it then follows by simple set-theoretic reasons that $(\ell, \ell') \in \mathbb{E}_2$ as required. Again, a symmetric chain of reasoning demonstrates that $\mathbb{E}_2 \subseteq \mathbb{E}_1$, hence we conclude.

$(\text{dom}(\rho_1) = \text{dom}(\rho_2))$: To see $\text{dom}(\rho_1) \subseteq \text{dom}(\rho_2)$, take $\ell \in \text{dom}(\rho_1)$. Since $\Sigma_1 \subseteq \Sigma'_1$, we have $\rho_1 \subseteq \rho'_1$ and thus that $\ell \in \text{dom}(\rho'_1)$. Then, since $\text{dom}(\rho'_1) = \text{dom}(\rho'_2)$, it follows that $\ell \in \text{dom}(\rho'_2)$. Also, $\ell \notin L$ by clause (3) of definition 22 since L is fresh for Σ_1 . Thus, since Σ'_2 is relevant for L over Σ_2 , we have by clause (2) of definition 21 that $\ell \notin \text{dom}(\rho'_2 \setminus \rho_2)$. However, since we have that $\ell \in \text{dom}(\rho'_2)$, by proposition 8 it must be that $\ell \in \text{dom}(\rho_2)$ as required. A symmetric chain of reasoning shows that $\text{dom}(\rho_2) \subseteq \text{dom}(\rho_1)$, hence we conclude.

$(\rho_1(\ell) \in \mathcal{V} \Leftrightarrow \rho_2(\ell) \in \mathcal{V})$: Assume $(\ell, v) \in \rho_1$ for some $v \in \mathcal{V}$; we show that there is some $v' \in \mathcal{V}$ with $(\ell, v') \in \rho_2$.

Since $\Sigma_1 \subseteq \Sigma'_1$, we have $\rho_1 \subseteq \rho'_1$ and thus that $(\ell, v) \in \rho'_1$. Then, since $\rho'_1(\ell) \in \mathcal{V} \Leftrightarrow \rho'_2(\ell) \in \mathcal{V}$, it follows that there is some $v' \in \mathcal{V}$ such that $(\ell, v') \in \rho'_2$. Also, $\ell \notin L$ since L is fresh for Σ_1 . Therefore, since Σ'_2 is relevant for L over Σ_2 , we have that $\ell \notin \text{dom}(\rho'_2 \setminus \rho_2)$. However, since we have that $(\ell, v') \in \rho'_2$, by proposition 8 it must be that $(\ell, v') \in \rho_2$ as required. A symmetric chain of reasoning shows that the converse direction holds, hence we conclude.

($\rho_1(\ell) = \rho_2(\ell)$ if $\rho_1(\ell) \notin \mathcal{V}$): Assume $\rho_1(\ell) = \iota$ and $\iota \notin \mathcal{V}$. Since $\Sigma_1 \subseteq \Sigma'_1$, we have $\rho_1 \subseteq \rho'_1$ and thus that $\rho'_1(\ell) = \iota$. Then, since $\rho'_1(\ell) \notin \mathcal{V}$ implies $\rho'_1(\ell) = \rho'_2(\ell)$, it follows that $\rho'_2(\ell) = \iota$. Also, $\ell \notin L$ since $\ell \in \text{dom}(\rho_1)$ (by assumption) and L is fresh for Σ_1 . Thus, since Σ'_2 is relevant for L over Σ_2 , we have by clause (2) of definition 21 that $\ell \notin \text{dom}(\rho'_2 \setminus \rho_2)$. However, since we have that $\rho'_2(\ell) = \iota$, by proposition 8 it must be that $\rho_2(\ell) = \iota$ as required. \square

This is used in the proofs of conjectures 14 and 15 and theorem 20 in order to infer the necessary conditions for applying the inductive hypothesis, namely relatedness of the semantics for corresponding sub-fragments of programs.

We now show some simple properties to do with preservation of properness and equivalence of semantics.

Lemma 11. *Suppose $\{\ell\}$ is fresh for Σ , with $\ell \neq \perp$; then Σ is proper if and only if $\Sigma[\ell \mapsto v]$ is.*

Proof. Immediate, by definition 19, since the only difference between the two semantics is the mapping of ℓ to v in the reification functions, and the freshness constraint entails that ℓ does not occur in the binding resolution function or the extension. \square

Lemma 12. *Let $v, v' \in \mathcal{V}$ with $\ell \notin \text{dom}(\Sigma_\rho)$, $\ell \notin \text{dom}(\Sigma'_\rho)$, $\ell \notin \text{ran}(\Gamma_1)$, and $\ell \notin \text{ran}(\Gamma_2)$ for $\ell \neq \perp$; then:*

1. $\Sigma \sim \Sigma'$ if and only if $\Sigma[\ell \mapsto v] \sim \Sigma'[\ell \mapsto v']$.
2. $\Gamma_1 \sim \Gamma_2$ only if $\Gamma_1[v \mapsto \ell] \sim \Gamma_2[v' \mapsto \ell]$.

Proof. Immediate, by definition 16. For the case of semantics, the result obtains because we have only updated the reification functions with mappings to value identifiers in both cases. For environments, we have only updated the value identifier mappings, in each case to the same location thus preserving the equality of the ranges. \square

Lemma 13. *let Σ and Σ' be semantics and ℓ a location such that there is no ℓ' such that $(\ell, \ell') \in \Sigma_{\mathbb{B}}$ or $(\ell, \ell') \in \Sigma'_{\mathbb{B}}$; then $\Sigma[\{\ell\} \otimes \Delta] \sim \Sigma'[\{\ell\} \otimes \Delta']$ implies $\Sigma \sim \Sigma'$, for all Δ, Δ' .*

Proof. The reification and binding resolution functions are not updated by the join operation. Thus it remains to show that $\Sigma_{\mathbb{B}} = \Sigma'_{\mathbb{B}}$. We show one direction of the inclusion; the other is symmetric. Let \mathbb{E}_+ and \mathbb{E}'_+ be the extension kernels of $\Sigma[\{\ell\} \otimes \Delta]$ and $\Sigma'[\{\ell\} \otimes \Delta']$, respectively. Suppose $(\ell_1, \ell_2) \in \Sigma_{\mathbb{B}}$. Since $\mathbb{E}_+ = \Sigma_{\mathbb{B}} \cup (\{\ell\} \otimes_{\Sigma_\rho} \Delta)$, thus also $(\ell_1, \ell_2) \in \mathbb{E}_+$. Since $\Sigma[\{\ell\} \otimes \Delta] \sim \Sigma'[\{\ell\} \otimes \Delta']$, it follows that $\mathbb{E}_+ = \mathbb{E}'_+$. Therefore $(\ell_1, \ell_2) \in \mathbb{E}'_+$. Notice that $\ell_1 \neq \ell$

since there is no ℓ' such that $(\ell, \ell') \in \Sigma_{\mathbb{B}}$. Moreover, by definition 11, all pairs in $\{\ell\} \otimes_{\Sigma_\rho} \Delta'$ are of the form (ℓ, ℓ') for some ℓ' . Thus $(\ell_1, \ell_2) \notin \{\ell\} \otimes_{\Sigma_\rho} \Delta'$. Since $\mathbb{E}'_+ = \Sigma'_{\mathbb{B}} \cup (\{\ell\} \otimes_{\Sigma_\rho} \Delta')$ it follows that we must have $(\ell_1, \ell_2) \in \Sigma'_{\mathbb{B}}$. \square

We now turn attention to the results of the renaming theory. Conjecture 1 is a corollary of the following property that we conjecture holds of our semantics. It should be possible to prove by induction on syntactic structure.

Conjecture 14. *If $\Sigma; \Gamma \vdash \sigma \hookrightarrow \sigma'$, with $\llbracket \sigma \rrbracket_{\Sigma; \Gamma} = (\mapsto', \mathbb{E}', \rho')$, then $\varphi(\sigma, \sigma') = U \cup L \cup C$, where:*

- $U \subseteq \{\ell \mid \ell \mapsto' \perp\}$,
- $L = \{\ell \mid \ell \in \delta(\sigma, \sigma') \vee \exists \ell' \in \delta(\sigma, \sigma'). \ell \mapsto' \ell'\}$, and
- $C \subseteq \{\ell \mid \exists \ell' \neq \perp. \ell' \in \text{decl}(\Sigma) \wedge \ell \mapsto' \ell'\}$.

From this we can immediately derive conjecture 1 by straightforwardly instantiating it with $\sigma \equiv P$ and $\sigma' \equiv P'$, and interpreting with respect to $\Sigma = \Sigma_\perp$ and $\Gamma = \Gamma_\perp$. In this case, notice that $C = \emptyset$.

Conjecture 1. *Suppose $P \hookrightarrow P'$ is a valid renaming, and let $L = \{\ell \mid \ell \in \delta(P, P') \vee \exists \ell' \in \delta(P, P'). \ell \mapsto_P \ell'\}$; then $L \subseteq \varphi(P, P')$ and $\ell \mapsto_P \perp$ for all $\ell \in \varphi(P, P') \setminus L$.*

Conjecture 2 is a corollary of the following property that we conjecture to hold of our semantics. Again, it should be possible to prove by induction on syntactic structure.

Conjecture 15. *Let $\Sigma_1 = (\mapsto_1, \mathbb{E}_1, \rho_1)$, $\Sigma_2 = (\mapsto_2, \mathbb{E}_2, \rho_2)$, such that both $\llbracket \sigma \rrbracket_{\Sigma_1; \Gamma_1}$ and $\llbracket \sigma' \rrbracket_{\Sigma_2; \Gamma_2}$ are defined and, moreover, $\llbracket \sigma \rrbracket_{\Sigma_1; \Gamma_1} \sim \llbracket \sigma' \rrbracket_{\Sigma_2; \Gamma_2}$; if D has a partitioning $P \subseteq \mathcal{L}_{\mathbb{B}^+}$, where $D = \{\ell \mid \ell \in (\text{dom}_{\mathcal{V}}(\rho_1) \setminus \text{dom}(\mapsto_1)) \wedge \rho_1(\ell) \neq \rho_2(\ell)\}$, then also $D \cup \delta(\sigma, \sigma')$ has a partitioning $P' \subseteq \mathcal{L}_{\mathbb{B}^+}$, where $\llbracket \sigma \rrbracket_{\Sigma_1; \Gamma_1} = (\mapsto', \mathbb{E}', \rho')$.*

Deriving conjecture 2 from this is done by straightforwardly instantiating it with $\sigma \equiv P$ and $\sigma' \equiv P'$, and interpreting with respect to $\Sigma_1 = \Sigma_2 = \Sigma_\perp$ and $\Gamma_1 = \Gamma_2 = \Gamma_\perp$; in this case, notice that $D = \emptyset$.

Conjecture 2. *Let $P \hookrightarrow P'$ be a valid renaming, then $\delta(P, P')$ has a partitioning that is a subset of $\mathcal{L}_{\mathbb{B}^+}$.*

Proposition 5 is a corollary of the following result.

Lemma 16. *Let $\llbracket \sigma \rrbracket_{\Sigma; \Gamma} \subseteq \Sigma'$, where for $\Sigma = (\mapsto_\Sigma, \mathbb{E}_\Sigma, \rho_\Sigma)$ and $\Sigma' = (\mapsto_{\Sigma'}, \mathbb{E}_{\Sigma'}, \rho_{\Sigma'})$, with Σ and Σ' proper; then for some given $\ell \in \text{decl}(\Sigma')$ and $v \in \mathcal{V}$ not occurring in σ or Σ' , define the following:*

- $L = \{\ell' \mid \ell' \in [\ell]_{\mathbb{B}} \vee \exists \ell'' \in [\ell]_{\mathbb{B}}. \ell' \mapsto \ell''\}$;
- $\sigma' = \sigma[\ell' \mapsto v \mid \ell' \in L \cap \text{dom}(\sigma)]$; and
- $\Sigma' = (\mapsto_{\Sigma'}, \mathbb{E}_{\Sigma'}, \rho_{\Sigma'}[\ell' \mapsto v \mid \ell' \in L \cap \text{dom}(\rho_{\Sigma'})])$.

Furthermore, define Γ' as follows: if there is (a necessarily unique) v' such that $\Gamma(v') \in [\ell]_{\mathbb{B}}$ then Γ' behaves as Γ except $\Gamma'(v) = \Gamma(v')$ and $\Gamma'(v') = \Gamma(v)$; otherwise, $\Gamma' = \Gamma$. Then $(\Sigma, \Gamma) \sim (\Sigma', \Gamma')$, $\llbracket \sigma' \rrbracket_{\Sigma'; \Gamma'}$ is defined, and $\llbracket \sigma \rrbracket_{\Sigma; \Gamma} \sim \llbracket \sigma' \rrbracket_{\Sigma'; \Gamma'}$.

Proof. By induction on syntactic structure. Given in the Coq formalisation. \square

Proposition 5. Suppose $\llbracket P \rrbracket$ is defined, $\ell \in \text{decl}(P)$, and $v \in \mathcal{V}$ does not occur in P , then $P \hookrightarrow P'$ is a valid renaming, where $P' = P[\ell' \mapsto v \mid \ell' \in [\ell]_{\mathbb{B}_P} \vee \exists \ell'' \in [\ell]_{\mathbb{B}_P}. \ell' \mapsto_P \ell'']$.

Proof. By straightforward instantiation of lemma 16 with $\sigma \equiv P$, interpreted with respect to $\Sigma = \Sigma_{\perp}$ and $\Gamma = \Gamma_{\perp}$. In this case, the definition of P' arises because we have by lemma 9 that $\llbracket P \rrbracket$ is relevant for $\text{dom}(P)$ over Σ_{\perp} and thus, by definition 21, it follows that $L \subseteq \text{dom}(P)$. \square

A.2 Adequacy

Here we give the full definition of our denotational model of behaviour for the OCaml module calculus. We first reiterate the definition of the denotational domain in which we interpret programs.

We assume an interpretation, using standard results, of value expressions (viz. lambda terms) in some domain \mathbb{F} containing an element **wrong** denoting run-time errors. We interpret modules in a domain \mathbb{M} satisfying:

$$\begin{aligned} \mathbb{M} &= \mathbb{D} + (\mathbb{M} \rightarrow \mathbb{M}) + \mathbf{wrong} \\ \mathbb{D} &= (\mathcal{V} \rightarrow_{\text{fin}} \mathbb{F}) \times (\mathcal{T} \rightarrow_{\text{fin}} \mathbb{T}) \times (\mathcal{M} \rightarrow_{\text{fin}} \mathbb{M}) \end{aligned}$$

where \mathbb{T} is the domain (defined below) in which we interpret module types. For $d \in \mathbb{D}$ we will write $\iota \in \text{dom}(d)$ to mean that ι is in the domain of the appropriate component of d , and $d(\iota)$ to mean the application of the appropriate component of d to ι . For $d, d' \in \mathbb{D}$, we write $d + d'$ for the module denotation (also in \mathbb{D}) for which $(d + d')(\iota) = d'(\iota)$ if $\iota \in \text{dom}(d')$, $(d + d')(\iota) = d(\iota)$ if $\iota \in \text{dom}(d) \setminus \text{dom}(d')$, and undefined otherwise. We define $d + \mathbf{wrong}$ to denote **wrong**. We will also sometimes describe an element $d \in \mathbb{D}$ as a (finite) set of pairs of the appropriate sorts of elements.

We interpret module types as elements in the initial algebra \mathbb{T} of the following functor F (in the category of sets):

$$\begin{aligned} F(X) &= D(X) + (\mathcal{M} \times X) \times X + \mathbf{wrong} \\ D(X) &= \wp_{\text{fin}}(\mathcal{V}) \times (\mathcal{T} \rightarrow_{\text{fin}} X) \times (\mathcal{M} \rightarrow_{\text{fin}} X) \end{aligned}$$

For $\tau \in D(\mathbb{T})$ we abuse notation and write $\mathcal{V}(\tau)$ for the first component of τ ; we also write $\tau(\iota)$ to mean the application of the appropriate component of τ to ι , and $\text{dom}(\tau)$ to mean the combined domains of the second and third components of τ . For $\tau, \tau' \in D(\mathbb{T})$ we write $\tau + \tau'$ for the module type denotation (also in $D(\mathbb{T})$) for which $\text{dom}_{\mathcal{V}}(\tau + \tau') = \text{dom}_{\mathcal{V}}(\tau) \cup \text{dom}_{\mathcal{V}}(\tau')$, with $(\tau + \tau')(\iota) = \tau'(\iota)$ if $\iota \in \text{dom}(\tau')$, $(\tau + \tau')(\iota) = \tau(\iota)$ if $\iota \in \text{dom}(\tau) \setminus \text{dom}(\tau')$, and undefined otherwise. We define $\tau + \mathbf{wrong}$ to denote **wrong**. We will also sometimes describe an element $\tau \in D(\mathbb{T})$ as a (finite) set of value identifiers and pairs of appropriate elements.

The denotational interpretation function $\llbracket \cdot \rrbracket_{\theta}$ is defined in fig. 4. It is parameterised by a denotational environment θ mapping value identifiers to elements of \mathbb{F} , module type identifiers to elements of \mathbb{T} , and module identifiers to pairs consisting of an element of \mathbb{M} and an element of \mathbb{T} . This function interprets value expressions in \mathbb{F} , module types in

\mathbb{T} , and module expressions as a pair of an element in \mathbb{M} and an element in \mathbb{T} . Thus, for a module expression, $\llbracket \cdot \rrbracket_{\theta}$ also synthesizes (the meaning of) its corresponding module type. We write $\llbracket \sigma \rrbracket$ to mean $\llbracket \sigma \rrbracket_{\theta_{\perp}}$, where θ_{\perp} is the environment that maps value and module type identifiers to **wrong** and module identifiers to the pair (**wrong**, **wrong**).

As a notational convenience, for $d \in \mathbb{D}$ and $\tau \in D(\mathbb{T})$ we write $\theta + (d, \tau)$ to denote the environment θ updated by the mappings in d , with mappings of module identifiers in d augmented by the corresponding module types in τ . That is, if $x \in \text{dom}(d)$, then $(\theta + (d, \tau))(x) = (d(x), \tau')$ where $\tau' = \tau(x)$ if $x \in \text{dom}(\tau)$ and $\tau' = \mathbf{wrong}$ otherwise.

The following coercion operation is used to give meaning to functors and module type annotations.

Definition 23 (Denotational Coercion). The (infix) operator $(:)$, of type $\mathbb{M} \times \mathbb{T} \rightarrow \mathbb{M}$, is defined inductively on the structure of module type denotations as follows.

$$d:\tau = \begin{cases} V \cup M \cup T & \text{if } d \in \mathbb{D} \wedge \tau \in D(\mathbb{T}) \\ \lambda d'.(d(d':\tau_1)):\tau_2 & \text{if } d \in \mathbb{M} \rightarrow \mathbb{M} \wedge \tau = ((x, \tau_1), \tau_2) \\ \mathbf{wrong} & \text{otherwise} \end{cases}$$

where $V = \{(v, d(v)) \mid v \in \text{dom}(d) \wedge v \in \mathcal{V}(\tau)\}$

$$\cup \{(v, \mathbf{wrong}) \mid v \notin \text{dom}(d) \wedge v \in \mathcal{V}(\tau)\}$$

$$M = \{(x, d(x):\tau(x)) \mid x \in \text{dom}(d) \wedge x \in \text{dom}(\tau)\}$$

$$\cup \{(x, \mathbf{wrong}) \mid x \notin \text{dom}(d) \wedge x \in \text{dom}(\tau)\}$$

$$T = \{(t, \tau(t)) \mid t \in \text{dom}(\tau)\}$$

We also define an operation to ‘promote’ a module type denotation to a module denotation. This operation reifies the structure of the module type denotation, building constant-valued functors for (sub)modules having a functor type. It is used to define the meaning of module types in various cases.

Definition 24 (Promotion). We define $(\cdot)^* : \mathbb{T} \rightarrow \mathbb{M}$ by induction on the structure of module type denotations.

$$\mathbf{wrong}^* = \mathbf{wrong}$$

$$\tau^* = \{(v, \mathbf{wrong}) \mid v \in \mathcal{V}(\tau)\} \quad \text{if } \tau \in D(\mathbb{T})$$

$$\cup \{(t, \tau(t)) \mid t \in \text{dom}(\tau)\}$$

$$\cup \{(x, \tau(x)^*) \mid x \in \text{dom}(\tau)\}$$

$$((x, \tau_1), \tau_2)^* = \lambda_{\cdot}.\tau_2^*$$

To prove the adequacy result, we must define how the elements of the set-theoretic semantics of section 3 relate to those of the denotational semantics defined in section 5. We first consider how the meanings of module types in the two semantics are related.

Definition 25. The relation $\tau \models_{\rho} \Delta$, for a module type denotation $\tau \in \mathbb{T}$ and a semantic description $\Delta \in \mathcal{D}$ w.r.t. a reification function ρ , is defined inductively as follows.

$$(1) \mathbf{wrong} \models_{\rho} \Delta \text{ for all } \Delta, \rho.$$

$$(2) \tau \models_{\rho} D, \text{ for } \tau \in D(\mathbb{T}), \text{ if:}$$

1981				2036
1982				2037
1983	$\llbracket x \rrbracket_\theta = \theta(x)$		$\llbracket q \cdot x \rrbracket_\theta = \text{match } \llbracket q \rrbracket_\theta \text{ with}$	2038
1984	$\llbracket q_1(q_2) \rrbracket_\theta = \text{let } (d', _) = \llbracket q_2 \rrbracket_\theta \text{ in match } \llbracket q_1 \rrbracket_\theta \text{ with}$		$ (d \in \mathbb{D}, \tau \in D(\mathbb{T})) \text{ when } x \in \text{dom}(\tau) \rightarrow$	2039
1985	$ (d \in \mathbb{M} \rightarrow \mathbb{M}, ((x, \tau_1), \tau_2)) \rightarrow (d(d'), \tau_2)$		$\text{if } x \in \text{dom}(d) \text{ then } (d(x), \tau(x)) \text{ else } (\mathbf{wrong}, \tau(x))$	2040
1986	$ (d \in \mathbb{M} \rightarrow \mathbb{M}, _) \rightarrow (d(d'), \mathbf{wrong})$		$ (d \in \mathbb{D}, \tau \in D(\mathbb{T})) \text{ when } x \in \text{dom}(d) \rightarrow (d(x), \mathbf{wrong})$	2041
1987	$ _ \rightarrow (\mathbf{wrong}, \mathbf{wrong})$		$ _ \rightarrow (\mathbf{wrong}, \mathbf{wrong})$	2042
1988		(a) Semantics of (extended) module paths.		2043
1989	$\llbracket t \rrbracket_\theta = \theta(t)$		$\llbracket \varepsilon \rrbracket_\theta = \emptyset$	2044
1990	$\llbracket p \cdot t \rrbracket_\theta = \text{let } (_, \tau) = \llbracket p \rrbracket_\theta \text{ in if } \tau \in D(\mathbb{T}) \text{ and } t \in \text{dom}(\tau)$		$\llbracket \mathbf{val } v : _ ; ; s \rrbracket_\theta = \text{let } \tau = \llbracket s \rrbracket_{\theta[v \mapsto \mathbf{wrong}]} \text{ in } \{v\} + \tau$	2045
1991	$\text{then } \tau(t) \text{ else } \mathbf{wrong}$		$\llbracket \mathbf{module } x : M ; ; s \rrbracket_\theta = \text{let } \tau = \llbracket M \rrbracket_\theta \text{ in}$	2046
1992	$\llbracket \mathbf{functor } (x : M_1) \rightarrow M_2 \rrbracket_\theta = \text{let } \tau = \llbracket M_1 \rrbracket_\theta \text{ in}$		$\text{let } \tau' = \llbracket s \rrbracket_{\theta[x \mapsto (\tau^*, \tau)]} \text{ in } \{(x, \tau)\} + \tau'$	2047
1993	$\text{let } \tau' = \llbracket M_2 \rrbracket_{\theta[x \mapsto (\tau^*, \tau)]} \text{ in } ((x, \tau), \tau')$		$\llbracket \mathbf{module type } t ; ; s \rrbracket_\theta = \text{let } \tau = \llbracket s \rrbracket_{\theta[t \mapsto \emptyset]} \text{ in } \{(t, \emptyset)\} + \tau$	2048
1994	$\llbracket M \text{ with module } x = q \rrbracket_\theta = \text{let } (_, \tau') = \llbracket q \rrbracket_\theta \text{ in let } \tau = \llbracket M \rrbracket_\theta \text{ in}$		$\llbracket \mathbf{module type } t = M ; ; s \rrbracket_\theta = \text{let } \tau = \llbracket M \rrbracket_\theta \text{ in}$	2049
1995	$\text{if } \tau \in D(\mathbb{T}) \text{ then } \llbracket M \rrbracket_{\theta[x \mapsto \tau']} \text{ else } \mathbf{wrong}$		$\text{let } \tau' = \llbracket s \rrbracket_{\theta[t \mapsto \tau]} \text{ in } \{(t, \tau)\} + \tau'$	2050
1996	$\llbracket M \text{ with module } x := q \rrbracket_\theta = \text{let } \tau = \llbracket M \rrbracket_\theta \text{ in}$		$\llbracket \mathbf{include } M ; ; s \rrbracket_\theta = \text{let } \tau = \llbracket M \rrbracket_\theta \text{ in if } \tau \in D(\mathbb{T}) \text{ then}$	2051
1997	$\text{if } \tau \in D(\mathbb{T}) \text{ then } \tau \setminus x \text{ else } \mathbf{wrong}$		$\text{let } \tau' = \llbracket s \rrbracket_{\theta+(\tau^*, \tau)} \text{ in } \tau + \tau'$	2052
1998	$\llbracket \mathbf{sig } S \text{ end} \rrbracket_\theta = \llbracket S \rrbracket_\theta$		$\text{else } \mathbf{wrong}$	2053
1999		(b) Semantics of module types.		2054
2000				2055
2001				2056
2002	$\llbracket \mathbf{struct } s \text{ end} \rrbracket_\theta = \llbracket s \rrbracket_\theta$		$\llbracket \mathbf{let } v = e ; ; s \rrbracket_\theta = \text{let } f = \llbracket e \rrbracket_\theta \text{ in let } (d, \tau) = \llbracket s \rrbracket_{\theta[v \mapsto f]} \text{ in}$	2057
2003	$\llbracket \mathbf{functor } (x : M) \rightarrow m \rrbracket_\theta = \lambda d. \text{if } d = \mathbf{wrong} \text{ then } \mathbf{wrong} \text{ else}$		$\{(v, f)\} + d, \{v\} + \tau$	2058
2004	$\text{let } \tau = \llbracket M \rrbracket_\theta \text{ in } \llbracket m \rrbracket_{\theta[x \mapsto (d, \tau)]}$		$\llbracket \mathbf{module } x = m ; ; s \rrbracket_\theta = \text{let } (d, \tau) = \llbracket m \rrbracket_\theta \text{ in}$	2059
2005	$\llbracket m_1(m_2) \rrbracket_\theta = \text{let } (d', _) = \llbracket m_2 \rrbracket_\theta \text{ in match } \llbracket m_1 \rrbracket_\theta \text{ with}$		$\text{let } (d', \tau') = \llbracket s \rrbracket_{\theta[x \mapsto (d, \tau)]} \text{ in}$	2060
2006	$ (d \in \mathbb{M} \rightarrow \mathbb{M}, ((x, \tau_1), \tau_2)) \rightarrow (d(d'), \tau_2)$		$\{(x, d)\} + d', \{(x, \tau)\} + \tau'$	2061
2007	$ (d \in \mathbb{M} \rightarrow \mathbb{M}, _) \rightarrow (d(d'), \mathbf{wrong})$		$\llbracket \mathbf{module type } t = M ; ; s \rrbracket_\theta = \text{let } \tau = \llbracket M \rrbracket_\theta \text{ in let } (d, \tau') = \llbracket s \rrbracket_{\theta[t \mapsto \tau]} \text{ in}$	2062
2008	$ _ \rightarrow (\mathbf{wrong}, \mathbf{wrong})$		$\{(t, \tau)\} + d, \{(t, \tau)\} + \tau'$	2063
2009	$\llbracket m : M \rrbracket_\theta = \text{let } d = \llbracket m \rrbracket_\theta \text{ in let } \tau = \llbracket M \rrbracket_\theta \text{ in } (d, \tau)$		$\llbracket \mathbf{include } m ; ; s \rrbracket_\theta = \text{let } (d, \tau) = \llbracket m \rrbracket_\theta \text{ in if } d \in \mathbb{D} \text{ then}$	2064
2010	$\llbracket \varepsilon \rrbracket_\theta = (\emptyset, \emptyset)$		$\text{let } (d', \tau') = \llbracket s \rrbracket_{\theta+(d, \tau)} \text{ in } (d + d', \tau + \tau')$	2065
2011			$\text{else } (\mathbf{wrong}, \mathbf{wrong})$	2066
2012		(c) Semantics of module expressions.		2067
2013				2068
2014				2069
2015	$\llbracket \mathbf{module } x = m ; ; P \rrbracket_\theta = \text{let } (d, \tau) = \llbracket m \rrbracket_\theta \text{ in } \llbracket P \rrbracket_{\theta[x \mapsto (d, \tau)]}$		$\llbracket p \cdot v \rrbracket_\theta = \text{let } (d, _) = \llbracket p \rrbracket_\theta \text{ in if } v \in \text{dom}(d) \text{ then } d(v) \text{ else } \mathbf{wrong}$	2070
2016		(d) Semantics of programs and module paths in value expressions.		2071
2017				2072
2018				2073
2019				2074
2020	(i) $\forall \ell \in D: \ell \in \text{dom}(\rho)$ and $\rho(\ell) \in \mathcal{V}(\tau)$; and		that two module denotations both constitute the same ‘imple-	2075
2021	(ii) $\forall (\ell, \Delta) \in D: \ell \in \text{dom}(\rho)$, $\rho(\ell) \in \text{dom}(\tau)$ and		mentation’ of a module description in the set-theoretic	2076
2022	$\tau(\rho(\ell)) \models_\rho \Delta$.		semantics with respect to two given reification functions.	2077
2023	(3) $((x, \tau), \tau') \models_\rho ((\ell, \Delta), \Delta')$ if:		Definition 26. For $\Delta \in \mathcal{D}$, $d, d' \in \mathbb{M}$, and reification func-	2078
2024	$\ell \in \text{dom}(\rho)$, $\rho(\ell) = x$, $\tau \models_\rho \Delta$, and $\tau' \models_\rho \Delta'$.		tions ρ, ρ' , the logical relation $\Delta \vdash (\rho, d) \sim (\rho', d')$ is defined	2079
2025	When $\tau \models_\rho \Delta$ holds, we say that the module type denotation		inductively on the structure of descriptions as follows.	2080
2026	τ models the semantic description Δ (w.r.t. ρ).			2081
2027	This relation satisfies a monotonicity property.			2082
2028	Lemma 17. <i>If $\tau \models_\rho \Delta$ and $\rho \subseteq \rho'$ then $\tau \models_{\rho'} \Delta$.</i>			2083
2029	<i>Proof.</i> Straightforward induction on the definition of \models_ρ . \square			2084
2030				2085
2031				2086
2032				2087
2033				2088
2034				2089
2035				2090

Figure 4. The denotational semantics of the OCaml calculus.

- 2091 iii. $\rho(\ell) \in \mathcal{M} \wedge \rho'(\ell) \in \mathcal{M} \Rightarrow$
 2092 $\Delta \vdash (\rho, d(\rho(\ell))) \sim (\rho', d'(\rho'(\ell)))$
 2093 3. $(\ell: \Delta_1 \rightarrow \Delta_2 \vdash (\rho, d) \sim (\rho', d'), \text{ for } d, d' \in \mathbb{M} \rightarrow \mathbb{M}, \text{ if:}$
 2094 $\forall d_1, d_2 \in \mathbb{M}, \rho_1 \supseteq \rho, \rho_2 \supseteq \rho'$
 2095 $\Delta_1 \vdash (\rho_1, d_1) \sim (\rho_2, d_2) \Rightarrow$
 2096 $\Delta_2 \vdash (\rho_1, d(d_1)) \sim (\rho_2, d'(d_2))$

2097 This logical relation is also monotone with respect to reifi-
 2098 cation functions.

2100 **Lemma 18.** *Suppose $\Delta \vdash (\rho_1, d_1) \sim (\rho_2, d_2)$, with $\rho_1 \vdash \Delta$ and*
 2101 *$\rho_2 \vdash \Delta$; if $\rho_1 \subseteq \rho'_1$ and $\rho_2 \subseteq \rho'_2$ then $\Delta \vdash (\rho'_1, d_1) \sim (\rho'_2, d_2)$.*

2102 *Proof.* Straightforward, by induction. \square

2104 Using these relations we can define when two sets of
 2105 semantics with a corresponding (set-theoretic) semantic and
 2106 denotational environment, constitute the same context up
 2107 to renaming.

2108 **Definition 27.** We define a relation on tuples of semantics
 2109 and environments by $(\Sigma, \Gamma, \theta) \sim (\Sigma', \Gamma', \theta')$ if and only if:

- 2110 (1) $(\Sigma, \Gamma) \sim (\Sigma', \Gamma')$;
 2111 (2) for all $v \in \mathcal{V}$,
 2112 (i) $\Gamma(v) = \perp \Rightarrow \theta(v) = \mathbf{wrong}$
 2113 (ii) $\Gamma'(v) = \perp \Rightarrow \theta'(v) = \mathbf{wrong}$
 2114 (iii) for all $v' \in \mathcal{V}$, $\Gamma(v) = \Gamma'(v') \Rightarrow \theta(v) = \theta'(v')$
 2115 (3) for all $t \in \mathcal{T}$, $\theta(t) \models_{\Sigma, \rho} \Gamma(t)$ and $\theta'(t) \models_{\Sigma', \rho'} \Gamma'(t)$
 2116 (4) for all $x \in \mathcal{M}$ with $\theta(x) = (d, \tau)$ and $\theta'(x) = (d', \tau')$,
 2117 (i) $\tau \models_{\Sigma, \rho} \Gamma(x)$ and $\tau' \models_{\Sigma', \rho'} \Gamma'(x)$
 2118 (ii) $\Delta \vdash (\Sigma, \rho, d) \sim (\Sigma', \rho', d')$, where $\Delta = \Gamma(x) = \Gamma'(x)$

2119 The following property holds.

2120 **Lemma 19.** *Let $v, v' \in \mathcal{V}$ be value identifiers, $d \in \mathbb{F}$ a value*
 2121 *denotation, and $\ell \neq \perp$ a location such that $\ell \notin \text{dom}(\Sigma, \rho)$,*
 2122 *$\ell \notin \text{dom}(\Sigma', \rho')$, $\ell \notin \text{ran}(\Gamma_1)$, and $\ell \notin \text{ran}(\Gamma_2)$ for $\ell \neq \perp$; then*

$$2123 \quad (\Sigma, \Gamma, \theta) \sim (\Sigma', \Gamma', \theta') \Rightarrow$$

$$2124 \quad (\Sigma[\ell \mapsto v], \Gamma[v \mapsto \ell], \theta[v \mapsto d]) \sim$$

$$2125 \quad (\Sigma'[\ell \mapsto v'], \Gamma'[v' \mapsto \ell], \theta'[v' \mapsto d])$$

2126 *Proof.* Suppose $(\Sigma, \Gamma, \theta) \sim (\Sigma', \Gamma', \theta')$. By lemma 12 it follows
 2127 that $\Sigma[\ell \mapsto v] \sim \Sigma'[\ell \mapsto v']$ and $\Gamma[v \mapsto \ell] \sim \Gamma'[v' \mapsto \ell]$. By
 2128 definition 27, the only extra condition we need to check is
 2129 clause 2(iii) for v and v' , since the additional mapping in
 2130 each case gives $\Gamma[v \mapsto \ell](v) = \Gamma'[v' \mapsto \ell](v') = \ell$. Notice
 2131 that we have $\theta[v \mapsto d](v) = \theta'[v' \mapsto d](v') = d$, and so the
 2132 condition is met. \square

2133 We can now show that the set-theoretic semantics refines
 2134 the denotational semantics.

2135 **Theorem 20 (Refinement).** *Suppose $\sigma_1 \hookrightarrow \sigma_2$ is a renaming,*
 2136 *$\llbracket \sigma_1 \rrbracket_{\Sigma_1; \Gamma_1} = \Sigma' \sim \Sigma'' = \llbracket \sigma_2 \rrbracket_{\Sigma_2; \Gamma_2}$ with Σ_1 and Σ_2 proper, and*
 2137 *there are θ_1, θ_2 with $(\Sigma_1, \Gamma_1, \theta_1) \sim (\Sigma_2, \Gamma_2, \theta_2)$; then:*

- 2138 1. *if σ_1, σ_2 are module types, then $\mathcal{D}_{\Sigma_1; \Gamma_1}(\sigma_1) = \mathcal{D}_{\Sigma_2; \Gamma_2}(\sigma_2) =$*
 2139 *Δ with $\llbracket \sigma_1 \rrbracket_{\theta_1} \models_{\Sigma', \rho} \Delta$ and $\llbracket \sigma_2 \rrbracket_{\theta_2} \models_{\Sigma'', \rho'} \Delta$;*

- 2140 2. *if σ_1, σ_2 are module expressions, where $\llbracket \sigma_1 \rrbracket_{\theta_1} = (d_1, \tau_1)$*
 2141 *and $\llbracket \sigma_2 \rrbracket_{\theta_2} = (d_2, \tau_2)$, then $\mathcal{D}_{\Sigma_1; \Gamma_1}(\sigma_1) = \mathcal{D}_{\Sigma_2; \Gamma_2}(\sigma_2) = \Delta$*
 2142 *with $\tau_1 \models_{\Sigma', \rho} \Delta$, $\tau_2 \models_{\Sigma'', \rho'} \Delta$, and $\Delta \vdash (\Sigma', d_1) \sim (\Sigma'', d_2)$;*
 2143 3. *if σ_1 and σ_2 are both value expressions or both programs,*
 2144 *then $\llbracket \sigma_1 \rrbracket_{\theta_1} = \llbracket \sigma_2 \rrbracket_{\theta_2}$.*

2145 *Proof.* By induction on syntactic structure. We show some
 2146 of the important cases in detail.

2147 **Value Expressions.** For value expressions, the result follows
 2148 straightforwardly by induction using the standard denota-
 2149 tional constructions of lambda calculus; we need only to
 2150 show that (qualified) value identifiers have the same denota-
 2151 tion. Let $\sigma \equiv p.v_\ell$ and $\sigma' \equiv p.v'_\ell$. Then $\Sigma' = \Sigma_3[\ell \mapsto (v, \ell')]$
 2152 and $\Sigma'' = \Sigma_4[\ell \mapsto (v', \ell')]$ for some ℓ' , where $\llbracket p \rrbracket_{\Sigma_1; \Gamma_1} = \Sigma_3 =$
 2153 $(\mapsto, \mathbb{E}, \rho)$ and $\llbracket p \rrbracket_{\Sigma_2; \Gamma_2} = \Sigma_4 = (\mapsto, \mathbb{E}, \rho)$, with $\Sigma_3 \sim \Sigma_4$. More-
 2154 over, by lemma 6, Σ_3 and Σ_4 are proper. Thus by the induc-
 2155 tive hypothesis we have that there is some $D = \mathcal{D}_{\Sigma_1; \Gamma_1}(p) =$
 2156 $\mathcal{D}_{\Sigma_2; \Gamma_2}(p)$ with $D \vdash (\rho, d_1) \sim (\rho', d_2)$, where $\llbracket p \rrbracket_{\theta_1} = (d_1, \tau_1)$
 2157 and $\llbracket p \rrbracket_{\theta_2} = (d_2, \tau_2)$. There are now two cases to consider,
 2158 from the definition of the set-theoretic semantics (cf. fig. 3):

2159 ($\ell' = \perp$): Then we have $\rho(\ell') \neq v$ and $\rho'(\ell') \neq v'$ for
 2160 all $\ell'' \in D$. Thus it follows from clauses (2a) and (2b) of
 2161 definition 26 that $v \notin \text{dom}(d_1)$ and $v' \notin \text{dom}(d_2)$. Therefore,
 2162 by definition (cf. fig. 4), $\llbracket p.v_\ell \rrbracket_{\theta_1} = \llbracket p.v'_\ell \rrbracket_{\theta_2} = \mathbf{wrong}$, as
 2163 required.

2164 ($\ell' \neq \perp$): Then we have that $\ell' \in D$ with $\rho(\ell') = v$ and
 2165 $\rho'(\ell') = v'$. It thus follows from clauses (2a) and (2b) of defini-
 2166 tion 26, respectively, that $v \in \text{dom}(d_1)$ and $v' \in \text{dom}(d_2)$, and
 2167 from clause (2c) that $d_1(v) = d_1(\rho(\ell')) = d_2(\rho'(\ell')) = d_2(v')$.
 2168 Therefore $\llbracket p.v_\ell \rrbracket_{\theta_1} = \llbracket p.v'_\ell \rrbracket_{\theta_2}$, as required.

2169 **Programs.** If σ_1 and σ_2 are value expressions, then the result
 2170 follows immediately from that for value expressions. When
 2171 $\sigma_1 \equiv \mathbf{let} \ x_\ell = m_1 \ ; \ ; \ ; \ P_1$ and $\sigma_2 \equiv \mathbf{let} \ x_\ell = m_2 \ ; \ ; \ ; \ P_2$, then
 2172 there are semantics $\Sigma_3 = \llbracket m_1 \rrbracket_{\Sigma_1; \Gamma_1}$ and $\Sigma_4 = \llbracket m_2 \rrbracket_{\Sigma_2; \Gamma_2}$ and
 2173 descriptions $\Delta_1 = \mathcal{D}_{\Sigma_1; \Gamma_1}(m_1)$ and $\Delta_2 = \mathcal{D}_{\Sigma_2; \Gamma_2}(m_2)$ such that
 2174 $\Sigma' = \llbracket P_1 \rrbracket_{\Sigma_3[\ell \mapsto x]; \Gamma_1[x \mapsto \Delta_1]}$ and $\Sigma'' = \llbracket P_2 \rrbracket_{\Sigma_4[\ell \mapsto x]; \Gamma_2[x \mapsto \Delta_2]}$. By
 2175 lemma 6, both Σ_3 and Σ_4 are proper. It thus follows trivially
 2176 from definition 19 that $\Sigma_3[\ell \mapsto x]$ and $\Sigma_4[\ell \mapsto x]$ are proper,
 2177 since the only difference in the updated semantics is in the
 2178 reification function. Therefore, by lemma 9, we have that
 2179 $\Sigma_3[\ell \mapsto x] \subseteq \Sigma'$ with $\Sigma' \setminus \Sigma_3[\ell \mapsto x] \subseteq \text{dom}(P_1)$ and $\text{dom}(P_1)$
 2180 fresh for $\Sigma_3[\ell \mapsto x]$, as well as $\Sigma_4[\ell \mapsto x] \subseteq \Sigma''$ with
 2181 $\Sigma'' \setminus \Sigma_4[\ell \mapsto x] \subseteq \text{dom}(P_2)$ and $\text{dom}(P_2)$ fresh for $\Sigma_4[\ell \mapsto x]$.
 2182 Hence, by lemma 10, $\Sigma_3[\ell \mapsto x] \sim \Sigma_4[\ell \mapsto x]$. More-
 2183 over notice that, by lemma 9, we have that $\Sigma_1 \subseteq \Sigma_3$ and
 2184 $\Sigma_3 \setminus \Sigma_1 \subseteq \text{dom}(m_1)$ with $\text{dom}(m_1)$ fresh for Σ_1 , and also
 2185 $\Sigma_2 \subseteq \Sigma_4$ and $\Sigma_4 \setminus \Sigma_2 \subseteq \text{dom}(m_2)$ with $\text{dom}(m_2)$ fresh for Σ_2 .
 2186 Therefore, given that neither $\ell \in \text{dom}(m_1)$ nor $\ell \in \text{dom}(m_2)$,
 2187 and $\Sigma_3[\ell \mapsto x] \sim \Sigma_4[\ell \mapsto x]$, it is then immediate from
 2188 definition 16 that $\Sigma_3 \sim \Sigma_4$. Now, since $\sigma_1 \hookrightarrow \sigma_2$ is a re-
 2189 naming, so is $m_1 \hookrightarrow m_2$. So, by the inductive hypothesis,
 2190 $\Delta_1 = \Delta_2 = \Delta$ with $\tau_1 \models_{\rho_3} \Delta$, $\tau_2 \models_{\rho_4} \Delta$ and $\Delta \vdash (\rho_3, d_1) \sim$
 2191 (ρ_4, d_2) , where $\llbracket m_1 \rrbracket_{\theta_1} = (\tau_1, d_1)$ and $\llbracket m_2 \rrbracket_{\theta_2} = (\tau_2, d_2)$, with

$\Sigma_3 = (\multimap_3, \mathbb{E}_3, \rho_3)$ and $\Sigma_4 = (\multimap_4, \mathbb{E}_4, \rho_4)$. Thus, according to definition 27, we straightforwardly obtain

$$\begin{aligned} & (\Sigma_3[\ell \mapsto x], \Gamma_1[x \mapsto \Delta], \theta_1[x \mapsto (d_1, \tau_1)]) \\ & \sim (\Sigma_4[\ell \mapsto x], \Gamma_2[x \mapsto \Delta], \theta_2[x \mapsto (d_2, \tau_2)]) \end{aligned}$$

since $(\Sigma_3[\ell \mapsto x], \Gamma_1[x \mapsto \Delta]) \sim (\Sigma_4[\ell \mapsto x], \Gamma_2[x \mapsto \Delta])$ is a precondition to the definedness of $\llbracket P_1 \rrbracket_{\Sigma_3[\ell \mapsto x], \Gamma_1[x \mapsto \Delta]}$ and $\llbracket P_2 \rrbracket_{\Sigma_4[\ell \mapsto x], \Gamma_2[x \mapsto \Delta]}$. Finally, this allows us to obtain from the inductive hypothesis that $\langle P_1 \rangle_{\theta_1[x \mapsto (d_1, \tau_1)]} = \langle P_2 \rangle_{\theta_2[x \mapsto (d_2, \tau_2)]}$, whence the result follows from the definition of the denotational semantics.

Value specifications. So $\sigma_1 \equiv \mathbf{val} \ v_\ell : _ ; ; S_1$ and $\sigma_2 \equiv \mathbf{val} \ v'_\ell : _ ; ; S_2$, with $\llbracket S_1 \rrbracket_{\Sigma'_1; \Gamma'_1} = \Sigma'_1$ and $\llbracket S_2 \rrbracket_{\Sigma'_2; \Gamma'_2} = \Sigma'_2$ where $\Sigma'_1 = \Sigma_1[\ell \mapsto v]$ and $\Sigma'_2 = \Sigma_2[\ell \mapsto v']$ with $\Gamma'_1 = \Gamma_1[v \mapsto \ell]$ and $\Gamma'_2 = \Gamma_2[v' \mapsto \ell]$. Then, let $D_1 = \mathcal{D}_{\Sigma'_1; \Gamma'_1}(S_1)$ and $D_2 = \mathcal{D}_{\Sigma'_2; \Gamma'_2}(S_2)$. So, $\Sigma' = \Sigma'_1[\{\ell\} \otimes D_1]$ and $\Sigma'' = \Sigma'_2[\{\ell\} \otimes D_2]$. Since Σ_1 and Σ_2 are proper, we have by lemma 11 that Σ'_1 and Σ'_2 are also proper. Moreover, since $\Sigma_1 \sim \Sigma_2$, we have by lemma 12(1) that $\Sigma'_1 \sim \Sigma'_2$. Similarly, since $\Gamma_1 \sim \Gamma_2$, we have by lemma 12(2) that $\Gamma'_1 \sim \Gamma'_2$. Now, take $\theta'_1 = \theta_1[v \mapsto \mathbf{wrong}]$ and $\theta'_2 = \theta_2[v' \mapsto \mathbf{wrong}]$. Thus, by lemma 19, we have $(\Sigma'_1, \Gamma'_1, \theta'_1) \sim (\Sigma'_2, \Gamma'_2, \theta'_2)$. Since $\ell \in \text{dom}(\sigma_1) = \text{dom}(\sigma_2)$, it follows from definition 15 and the fact that Γ_1 and Γ_2 are well-behaved w.r.t. Σ_1 and Σ_2 , respectively, that $\ell \notin \text{ran}_{\mathcal{D}}(\Gamma_1)$ and $\ell \notin \text{ran}_{\mathcal{D}}(\Gamma_2)$. Thus $\ell \notin \text{ran}_{\mathcal{D}}(\Gamma'_1)$ and $\ell \notin \text{ran}_{\mathcal{D}}(\Gamma'_2)$, also. Notice too that $\ell \notin \text{dom}(S_1) = \text{dom}(S_2)$. Therefore, by lemma 7, $\ell \notin D_1$ and $\ell \notin D_2$. Furthermore, since ℓ does not appear in Σ_1 or Σ_2 (by definition 15), it follows from lemma 8 that there is no ℓ' such that (ℓ, ℓ') is contained in the extension kernels of Σ'_1 or Σ'_2 . Thus by lemma 13 we have $\Sigma'_1 \sim \Sigma'_2$. So, by the inductive hypothesis, we obtain $D_1 = D_2 = D$ with $\langle S_1 \rangle_{\theta'_1} \models_{\rho'_1} D$ and $\langle S_2 \rangle_{\theta'_2} \models_{\rho'_2} D$, where ρ'_1 and ρ'_2 are the reification functions of Σ'_1 and Σ'_2 , respectively. Then $\mathcal{D}_{\Sigma_1; \Gamma_1}(\sigma_1) = \{\ell\} \oplus_{\Sigma'} D$ and $\mathcal{D}_{\Sigma_2; \Gamma_2}(\sigma_2) = \{\ell\} \oplus_{\Sigma''} D$. We also have, by lemma 9, that $\Sigma'_1 \subseteq \Sigma'_1$ and $\Sigma'_2 \subseteq \Sigma'_2$. Let $\Sigma'_1 = (\multimap', \mathbb{E}', \rho')$ and $\Sigma'_2 = (\multimap'', \mathbb{E}'', \rho'')$. Then

- $\Sigma' = (\multimap', \mathbb{E}' \cup (\{\ell\} \otimes_{\rho'} D), \rho')$; and
- $\Sigma'' = (\multimap'', \mathbb{E}'' \cup (\{\ell\} \otimes_{\rho''} D), \rho'')$.

We now need to show the following.

1. $\mathcal{D}_{\Sigma_1; \Gamma_1}(\sigma_1) = \mathcal{D}_{\Sigma_2; \Gamma_2}(\sigma_2)$, i.e. $\{\ell\} \oplus_{\rho'} D = \{\ell\} \oplus_{\rho''} D$. By definition 10, it suffices to prove $\exists \ell' \in D. \rho'(\ell) = \rho''(\ell')$ if and only if $\exists \ell'' \in D. \rho''(\ell'') = \rho'(\ell')$. We show the ‘only if’ direction; the other is symmetric. Assume $\ell' \in D$ with $\rho'(\ell) = \rho''(\ell')$. Then by definition 11 $(\ell, \ell') \in \{\ell\} \otimes_{\rho'} D$. Therefore, since $\Sigma' \sim \Sigma''$, we have by definition 16 that also $(\ell, \ell') \in \mathbb{E}'' \cup (\{\ell\} \otimes_{\rho''} D)$. The result is then obtained immediately from definition 19 since, by lemma 6, Σ'' is proper and so $\rho''(\ell) = \rho''(\ell')$.

2. $\langle S_1 \rangle_{\theta_1} \models_{\rho'} D'$ and $\langle S_2 \rangle_{\theta_2} \models_{\rho''} D'$, for $D' = \{\ell\} \oplus_{\rho'} D = \{\ell\} \oplus_{\rho''} D$. We show that $\langle S_1 \rangle_{\theta_1} \models_{\rho'} D'$; showing the other is similar. We distinguish two cases.

- If there exists some $\ell' \in D$ such that $\rho'(\ell') = v$ then $D' = \{\ell\} \oplus_{\rho'} D = D$ and, by clause 2(i) of definition 25, $\langle S_1 \rangle_{\theta_1} = \{v\} + \langle S_1 \rangle_{\theta'_1} = \langle S_1 \rangle_{\theta'_1}$ since $v \in \mathcal{V}(\langle S_1 \rangle_{\theta'_1})$. Therefore the result follows, by lemma 17, from the fact that $\langle S_1 \rangle_{\theta'_1} \models_{\rho'_1} D$ and $\rho'_1 \subseteq \rho'$, the latter entailed by $\Sigma'_1 \subseteq \Sigma'_1$.
- Otherwise, then $\ell \in D' = \{\ell\} \cup D$ and $v \in \mathcal{V}(\langle S_1 \rangle_{\theta_1}) = \{v\} \cup \langle S_1 \rangle_{\theta'_1}$. Since $\Sigma'_1 \subseteq \Sigma'_1$, and thus $\rho'_1 \subseteq \rho'$, we have by lemma 17 that $\langle S_1 \rangle_{\theta'_1} \models_{\rho'} D$. Notice that we also thus have $\rho'(\ell) = v$ since $\rho'_1(\ell) = v$. The result then follows straightforwardly by definition 25.

Value definitions. This is similar to the case for value specifications above. Here we have $\sigma_1 \equiv \mathbf{let} \ v_\ell = e_1 ; ; s_1$ and $\sigma_2 \equiv \mathbf{let} \ v'_\ell = e_2 ; ; s_2$ with $\llbracket e_1 \rrbracket_{\Sigma_1; \Gamma_1} = \Sigma_3$, $\llbracket e_2 \rrbracket_{\Sigma_2; \Gamma_2} = \Sigma_4$, $\llbracket s_1 \rrbracket_{\Sigma'_3; \Gamma'_1} = \Sigma'_3$, and $\llbracket s_2 \rrbracket_{\Sigma'_4; \Gamma'_2} = \Sigma'_4$ where $\Sigma'_3 = \Sigma_3[\ell \mapsto v]$, $\Gamma'_1 = \Gamma_1[v \mapsto \ell]$, $\Sigma'_4 = \Sigma_4[\ell \mapsto v']$, and $\Gamma'_2 = \Gamma_2[v' \mapsto \ell]$. Moreover, let $D'_1 = \mathcal{D}_{\Sigma'_3; \Gamma'_1}(s_1)$ and $D'_2 = \mathcal{D}_{\Sigma'_4; \Gamma'_2}(s_2)$. So, $\Sigma' = \Sigma'_3[\{\ell\} \otimes D'_1]$ and $\Sigma'' = \Sigma'_4[\{\ell\} \otimes D'_2]$. Since $\ell \in \text{dom}(\sigma_1) = \text{dom}(\sigma_2)$, it follows from definition 15 and the fact that Γ_1 and Γ_2 are well-behaved w.r.t. Σ_1 and Σ_2 , respectively, that $\ell \notin \text{ran}_{\mathcal{D}}(\Gamma_1)$ and $\ell \notin \text{ran}_{\mathcal{D}}(\Gamma_2)$. Thus $\ell \notin \text{ran}_{\mathcal{D}}(\Gamma'_1)$ and $\ell \notin \text{ran}_{\mathcal{D}}(\Gamma'_2)$, also. Notice too that $\ell \notin \text{dom}(s_1) = \text{dom}(s_2)$. Therefore, by lemma 7, $\ell \notin D_1$ and $\ell \notin D_2$. Furthermore, since ℓ does not appear in Σ_1 or Σ_2 (by definition 15), nor in $\text{dom}(e_1) = \text{dom}(e_2)$, it follows from lemma 9(2) that ℓ does not appear in Σ_3 or Σ_4 . Then, by lemma 8, we have that there is no ℓ' such that (ℓ, ℓ') is contained in the extension kernels of Σ'_3 or Σ'_4 . Thus by lemma 13 we have $\Sigma'_3 \sim \Sigma'_4$. Now, by lemma 9 we have that $\Sigma'_3 \subseteq \Sigma'_3$ and $\Sigma'_4 \subseteq \Sigma'_4$ (properly), as well as $\Sigma'_3 \setminus \Sigma'_3 \subseteq L$ and $\Sigma'_4 \setminus \Sigma'_4 \subseteq L$ with L fresh for both Σ'_3 and Σ'_4 , where $L = \text{dom}(s_1) = \text{dom}(s_2)$. So, by lemma 10, it follows that $\Sigma'_3 \sim \Sigma'_4$ and therefore, by lemma 12(1), that $\Sigma_3 \sim \Sigma_4$. Thus, by the inductive hypothesis, we have $\langle e_1 \rangle_{\theta_1} = \langle e_2 \rangle_{\theta_2} = d$. Furthermore, by lemma 6, both Σ_3 and Σ_4 are proper. Thus, by lemma 11 it follows that Σ'_3 and Σ'_4 are proper too. We also have, by lemma 12(2) that $\Gamma'_1 \sim \Gamma'_2$. Now, take $\theta'_1 = \theta_1[v \mapsto d]$ and $\theta'_2 = \theta_2[v' \mapsto d]$. It then follows from lemma 19 that $(\Sigma'_3, \Gamma'_1, \theta'_1) \sim (\Sigma'_4, \Gamma'_2, \theta'_2)$. Thus, another application of the inductive hypothesis derives that $D'_1 = D'_2 = D'$ with $\tau'_1 \models_{\rho'} D'$, $\tau'_2 \models_{\rho''} D'$, and $D' \vdash (\rho', d'_1) \sim (\rho'', d'_2)$, for $(d'_1, \tau'_1) = \langle S_1 \rangle_{\theta'_1}$ and $(d'_2, \tau'_2) = \langle S_2 \rangle_{\theta'_2}$, where $\Sigma'_3 = (\multimap', \mathbb{E}', \rho')$, and $\Sigma'_4 = (\multimap'', \mathbb{E}'', \rho'')$. We must now show three things.

- (i) $\mathcal{D}_{\Sigma_1; \Gamma_1}(\sigma_1) = \{\ell\} \otimes_{\rho'} D' = \{\ell\} \otimes_{\rho''} D' = \mathcal{D}_{\Sigma_2; \Gamma_2}(\sigma_2)$.
- (ii) $\{v\} + \tau'_1 \models_{\rho'} \{\ell\} \otimes_{\rho'} D'$ and $\{v'\} + \tau'_2 \models_{\rho''} \{\ell\} \otimes_{\rho''} D'$.
- (iii) $D \vdash (\rho', d_1) \sim (\rho'', d_2)$, where $\mathcal{D}_{\Sigma_1; \Gamma_1}(\sigma_1) = \mathcal{D}_{\Sigma_2; \Gamma_2}(\sigma_2) = D$ with $d_1 = \{(v, d)\} + d'_1$ and $d_2 = \{(v', d)\} + d'_2$.

The first two properties hold by the same reasoning as shown in the case for value descriptions above. To show that the last property holds, we consider two cases.

($D = D'$): So there is $\ell' \in D'$ such that $\rho'(\ell') = v$ and $\ell'' \in D'$ such that $\rho''(\ell'') = v'$. Thus, since $D' \vdash (\rho', d'_1) \sim (\rho'', d'_2)$ we have by definition 26 that $v \in \text{dom}(d'_1)$

2311 and $v' \in \text{dom}(d'_2)$. Therefore, $\{(v, d)\} + d'_1 = d'_1$ and 2366
 2312 $\{(v', d)\} + d'_2 = d'_2$, whence the result follows directly. 2367
 2313 ($\ell \notin D'$, $D = D' \cup \{\ell\}$): Since $D' \vdash (\rho', d'_1) \sim (\rho'', d'_2)$ it fol- 2368
 2314 lows from definition 26 that $v \notin \text{dom}(d'_1)$ and $v' \notin$ 2369
 2315 $\text{dom}(d'_2)$. Thus, we have that $\text{dom}(d_1) = \text{dom}(d'_1) \cup \{v\}$ 2370
 2316 and $\text{dom}(d_2) = \text{dom}(d'_2) \cup \{v'\}$. Moreover, $d_1(v) = d =$ 2371
 2317 $d_2(v')$. From these properties, we can derive the result 2372
 2318 by definition 26. \square 2373

2319 **Proposition 6** (Adequacy). $\langle\!\langle P \rangle\!\rangle = \langle\!\langle P' \rangle\!\rangle$ if $P \leftrightarrow P'$ is valid. 2374
 2320

2321 *Proof.* By straightforward instantiation of theorem 20 with 2375
 2322 $\sigma \equiv P$ and $\sigma' \equiv P'$, interpreted with respect to $\Sigma_1 = \Sigma_2 = \Sigma_\perp$, 2376
 2323 $\Gamma_1 = \Gamma_2 = \Gamma_\perp$, and $\theta_1 = \theta_2 = \theta_\perp$, for which it is straightfor- 2377
 2324 ward to show that $(\Sigma_\perp, \Gamma_\perp, \theta_\perp) \sim (\Sigma_\perp, \Gamma_\perp, \theta_\perp)$. \square 2378
 2325 2379
 2326 2380
 2327 2381
 2328 2382
 2329 2383
 2330 2384
 2331 2385
 2332 2386
 2333 2387
 2334 2388
 2335 2389
 2336 2390
 2337 2391
 2338 2392
 2339 2393
 2340 2394
 2341 2395
 2342 2396
 2343 2397
 2344 2398
 2345 2399
 2346 2400
 2347 2401
 2348 2402
 2349 2403
 2350 2404
 2351 2405
 2352 2406
 2353 2407
 2354 2408
 2355 2409
 2356 2410
 2357 2411
 2358 2412
 2359 2413
 2360 2414
 2361 2415
 2362 2416
 2363 2417
 2364 2418
 2365 2419
 2420