Multi-stage Programming with Splice Variables

TSUNG-JU CHIANG, University of Toronto, Canada
NINGNING XIE, University of Toronto, Canada

Multi-stage programming is a popular approach to typed meta-programming, reducing abstraction overhead and producing performant programs. However, the traditional quote-and-splice staging syntax, as introduced by Rowan Davies in 1996, can introduce complexities in managing expression evaluation, and also often necessitates sophisticated mechanisms for advanced features such as code pattern matching. This paper introduces $\lambda^{\bigcirc\triangleright}$, a novel staging calculus featuring let-splice bindings, a construct that explicitly binds splice expressions to splice variables, providing flexibility in managing, sharing, and reusing splice computations. Inspired by contextual modal type theory, our type system associates types with a typing context to capture variables dependencies of splice variables. We demonstrate that this mechanism seamlessly scales to features like code pattern matching, by formalizing $\lambda^{\bigcirc\triangleright}_{\text{pat}}$, an extension of $\lambda^{\bigcirc\triangleright}$ with code pattern matching and rewriting. We establish the syntactic type soundness of both calculi. Furthermore, we define a denotational semantics using a Kripke-style model, and prove adequacy results. All proofs have been fully mechanized using the Agda proof assistant.

CCS Concepts: • **Software and its engineering** → **Functional languages**; **Semantics**; • **Theory of computation** → **Type theory**.

Additional Key Words and Phrases: Staging, Temporal logic, Code Pattern Matching, Agda

ACM Reference Format:
Tsung-Ju Chiang and Ningning Xie. 2025. Multi-stage Programming with Splice Variables. *Proc. ACM Program. Lang.* 9, ICFP, Article 249 (August 2025), 35 pages. https://doi.org/10.1145/3747518

1 Introduction

Typed meta-programming allows programs to generate and manipulate code fragments, reducing abstraction overhead and producing performant programs. A prominent approach to typed meta-programming is *multi-stage programming* based on temporal logic [Davies 1996, 2017], which has been implemented in several languages, including OCaml [Chiang et al. 2024; Kiselyov 2014; Xie et al. 2023], Scala [Stucki et al. 2018, 2021], and Haskell [Sheard and Peyton Jones 2002; Xie et al. 2022], and has been successfully applied to improve performance in various domains [Jonnalagedda et al. 2014; Kiselyov et al. 2017; Wang et al. 2019; Willis et al. 2020; Yallop 2017; Yallop et al. 2023]. Multi-stage programming integrates type and scope checking of code expressions into the type system, allowing meta-programs to be type-checked at compile time, leading to safer and more predictable code generation.

A classic example of staged programming is the power function. We present the example below using the *quoting* and *splicing* mechanisms, where a quotation `<..>` represents the code fragment of an expression, and a splice `$(..)` extracts the term from a quoted code fragment:

```
let rec power x n = if n == 0 then <1>
                    else <$(x) * $(power x (n - 1))>
```

Authors' Contact Information: Tsung-Ju Chiang, University of Toronto, Toronto, Canada, mail@tsung-ju.org; Ningning Xie, University of Toronto, Toronto, Canada, ningningxie@cs.toronto.edu.

ACM 2475-1421/2025/8-ART249
https://doi.org/10.1145/3747518

Given a quoted expression `<e>` and a known integer `n`, this function generates an expression with n repeated multiplications, avoiding recursion at runtime and thus reducing overhead for any specific `e`. For example, with `n` being `5`, we have:

```
<fun x → $(power <x> 5)>    -- generates <fun x → x * x * x * x * x * 1>
```

While the quote-and-splice syntax offers a useful staging mechanism, it also introduces several complexities. Specifically, managing expression evaluation is subtle and often involved a *level-index reduction relation* [Calcagno et al. 2003; Taha and Sheard 2000], complicating both implementations and formal semantics. Moreover, supporting pattern matching on code often requires sophisticated mechanisms [Jang et al. 2022; Parreaux et al. 2017; Stucki et al. 2018].

In this paper we introduce *let-splice bindings*, a language construct that explicitly binds splice expressions to *splice variables*. Unlike the splice operator, let-splice bindings offer precise control over splice evaluation order, drawing inspiration from the target calculus of Typed Template Haskell [Xie et al. 2022]. Moreover, let-splice bindings provide greater flexibility and enable the sharing and reuse of splice computations. Our design includes a novel type system that tracks the *variable dependencies* of splice variables, allowing splice expressions to be defined in contexts where certain variables are not yet available. Inspired by *contextual modal type theory* [Jang et al. 2022; Nanevski et al. 2008], this type system associates types with a typing context to capture these variable dependencies, ensuring well-typed splice definitions. When a splice variable is used, its associated dependencies must be provided. These contexts can be nested to specify intricate dependency structures. Furthermore, we demonstrate that our design naturally supports advanced meta-programming features such as *unhygienic functions* [Barzilay et al. 2011], *code pattern matching* [Stucki et al. 2021], and *code rewriting* [Parreaux et al. 2017], using the same mechanism for managing variable dependencies, highlighting the expressiveness of our approach.

We present two calculi: (1) $\lambda^{\bigcirc\triangleright}$, a temporal-style multi-stage calculus supporting let-splice bindings, featuring a novel contextual modality ($\Delta \triangleright$) for managing variable dependencies; and (2) $\lambda^{\bigcirc\triangleright}_{\text{pat}}$, an extension of $\lambda^{\bigcirc\triangleright}$ which integrates code pattern matching and code rewriting. Both calculi are equipped with a small-step operational semantics, and enjoy syntactic type soundness. Moreover, we define a denotational semantics based on a Kripke-style model for both calculi, and establish the adequacy results for $\lambda^{\bigcirc\triangleright}$. The calculi along with all proofs are fully formalized in the Agda proof assistant. In particular, we offer the following contributions:

- §3 and 4 present $\lambda^{\bigcirc\triangleright}$, a staging calculus featuring let-splice bindings with splice variables, dependency tracking with nested typing context, a temporal-style code type for code expressions, and a separate contextual modality for managing variable dependencies.
- §5 proves soundness and completeness of $\lambda^{\bigcirc\triangleright}$'s type system with respect to *constructive linear-time temporal logic* [Kojima and Igarashi 2011], and then provides a type-preserving translation from $\lambda^{\bigcirc\triangleright}$ to $\lambda^{\bigcirc}$ [Davies 1996], providing insight into their relationship.
- §6 introduces $\lambda^{\bigcirc\triangleright}_{\text{pat}}$, an extension of $\lambda^{\bigcirc\triangleright}$ that enables pattern matching on code and code rewriting, allowing for inspection and manipulation of code expressions.
- §7 defines a denotational semantics for $\lambda^{\bigcirc\triangleright}$ and $\lambda^{\bigcirc\triangleright}_{\text{pat}}$ using a Kripke-style model, establishing termination of the operational semantics and adequacy of the denotational semantics.
- We formalize $\lambda^{\bigcirc\triangleright}$ and $\lambda^{\bigcirc\triangleright}_{\text{pat}}$ in the Agda proof assistant (§8), and establish key properties and theorems including progress, preservation, and adequacy. The Agda definitions and proofs of all stated lemmas and theorems are provided in the supplementary materials.

Lastly, §9 compares our approach to related work. For reasons of space, some of the rules of our calculi are elided in the paper; the complete set of rules are provided in the appendix.

2 Examples

This section outlines the design of our calculus and illustrates its expressiveness through examples. We use OCaml-like syntax throughout this section.

2.1 Staged Power Function

Recall the staged power function from the introduction:

```
val power : int¹ code → int⁰ → int¹ code
let rec power x n = if n == 0 then <1>
                    else <$(x) * $(power x (n - 1))>
```

Here, a *quotation* `<expr>` represents the code fragment of the expression, and a *splice* `$(expr)` extracts out the expression from the code fragment. The type constructor `code` annotates typed code expressions. Formally, the type system associates each variable and expression with a *level* (i.e. "time" in temporal logic). Intuitively, levels indicate the evaluation stage of expressions. Quotations increase the level of the quoted expression, splices decrease the level, and local variables have the same level as their surrounding context and can only be used at that level. For clarity, we annotate base types with their levels.

While the quote-and-splice syntax is useful, it also introduces several complexities. Consider evaluating `(e1 <e2 $(e3)>)`. Here, `e1` and `e3` are evaluated, but `e2` is not, as `e2` represents an expression in the next stage. Therefore, evaluating a staged program necessitates a *level-indexed reduction relation* [Calcagno et al. 2003; Taha and Sheard 2000] to track an evaluation level, which is dynamically updated by quotations and splices during evaluation. This adds complexity to both implementations and formal semantics. Moreover, it is difficult to precisely control the evaluation of splice expressions (see also §2.2). Such fine-grain control is also needed for *nested splices* in compile-time code generation systems [Sheard and Peyton Jones 2002; Xie et al. 2022, 2023], where splices without surrounding quotations are evaluated at compile-time. For instance, in `$(e1) $($(e2))`, the relative evaluation order of `e1` and `e2` is unclear: while `e1` appears first, `e2` has more splices. Recently, Xie et al. [2022] addresses this by prioritizing more deeply nested splices through an elaboration process. However, the resulting target calculus is not directly accessible to programmers.

In this paper, we present *let-splice bindings*, a new construct that defines *splice variables*. This approach explicitly defines splice expressions and thus specifies their evaluation order, eliminating the need for a level-indexed reduction relation. For example, the staged power function can be written as follows using our syntax:

```
val power : int¹ code → int⁰ → int¹ code
let rec power x n = if n == 0 then <1>
                    else let$ s1 : int¹ = x in                  -- lifted
                         let$ s2 : int¹ = power x (n - 1) in    -- lifted
                         <s1 * s2>
```

Here, let-splice bindings, defined using `let$`, replace the splice operation. There definitions bind a code expression to a splice variable. In this case, the splice variables `s1` and `s2` represent the splices of `x` and `power x (n - 1)`, respectively. Since these variables represent splices in the original program, they can be used directly within the quotation, as in `<s1 * s2>`. Formally, let-splice bindings unwrap a `code`-typed expression, introducing a variable which is one level higher and thus may be used within quotations. Let-splice bindings are evaluated sequentially; in this case, `s1` is evaluated before `s2`, making the evaluation order of splice expressions explicit.

More interestingly, let-splice bindings can be annotated with a list of *variable dependencies*. This provides flexibility for splice expressions which depend on values that are only available when the splice variable is used. For instance, we can define `power5` as:

```
let$ power5 : (x : int¹ ⊢ int¹)  -- variable dependency on x
   = power <x> 5                 -- generates <x * x * x * x * x * 1>
```

As `(power <x> 5)` refers to the variable `x`, the splice variable `power5` is declared with a dependency on `x : int¹`, allowing `x` to be used within its definition. Our type system tracks splice variables' dependencies, in addition to their levels. In this case, `power5` has type `(x : int¹ ⊢ int¹)`, indicating that it is a level-1 splice with type `int¹` that depends on a variable `x : int¹`.

To use a splice variable with variable dependencies like `power5`, the syntax `(s with x = e)` provides a *delayed substitution*. This allows variable dependencies to be replaced with specific values when a splice variable is actually used:

```
let thirty_two = <power5 with x = 2>  -- <2 * 2 * 2 * 2 * 2 * 1>
```

The right-hand side of the substitution can be any expression or variable in general. For brevity, we often write `s with x` as a shorthand for `s with x = x`. Similarly, `s with x; y` is a shorthand for `s with x = x; y = y`, and so on. For example:

```
let fpower = <fun x ⟶ power5 with x>  -- <fun x ⟶ x * x * x * x * x * 1>
```

In practice, an implementation may automatically capture dependencies from the context, allowing identical entries like `with x = x` to be omitted entirely.

2.2 Reuse of Splice Variables

We demonstrate that let-splice bindings and splice variables offer finer-grained control over splice expression evaluation. In particular, consider below a program using the quote-and-splice syntax, which generates a pair of functions that apply a function `f : int¹ code ⟶ int¹ code` to their arguments, and increment and decrement the result, respectively:

```
<(fun x ⟶ $(f <x>) + 1, fun y ⟶ $(f <y>) - 1)>  -- quote and splice
```

For example, given `f z = (let w = very_long_computation in <$(z) * 2>)`, this generates

```
<(fun x ⟶ x * 2 + 1, fun y ⟶ y * 2 - 1)>
```

However, the program evaluates the two splices `$(f <x>)` and `$(f <y>)` sequentially, leading to duplicated computations of `very_long_computation`. Similarly, it may also be desirable to run a function `f` only once when it is effectful.

To avoid duplicated computations, we may pre-compute the splice expression's result (here `z` is introduced and used at level 1), bind it to a variable, and use that variable in the definition:

```
let s = <fun z ⟶ $(f <z>)> in
<(fun x ⟶ $(s) x + 1, fun y ⟶ $(s) y - 1)>
```

Unfortunately, the generated code is different:

```
<(fun x ⟶ ((fun z ⟶ z * 2) x) + 1, fun y ⟶ ((fun z ⟶ z * 2) y) - 1)>
```

introducing two additional beta-redexes, since `s` produces a function being applied to `x`. This may not be desirable, for example, when we can inspect the generated code using pattern matching (§2.5).[1]

[1]One may also write `let s = fun z ⟶ <$(f z)> in <(fun x ⟶ $(s <x>) + 1, fun y ⟶ $(s <y>) - 1)>`, which however still results in `f` being evaluated twice.

In our calculus, we can easily share splice computations using splice variables. Specifically, we can write the original computation as:

```
let$ s : (z : int¹ ⊢ int¹) = f <z> in  -- splice variable
<(fun x → (s with z = x) + 1, fun y → (s with z = y) - 1)>
```

Here, `let$` declares a splice variable `s` with a dependency on `z`. The expression `f <z>` is evaluated, which can refer to variable `z`. The `(s with z = x)` syntax directly substitutes `z` with `x`, and `(s with z = y)` substitutes `z` with `y`. In this case, the `very_long_computation` is evaluated only once, and the generated code is the desired `<(fun x → x * 2 + 1, fun y → y * 2 - 1)>`.

2.3 Variable Dependencies

Building upon variable dependencies, our system extends support for dependency declarations to standard let bindings and function arguments. As an example, consider the following program:

```
val w : (x : string¹; y : int¹ ⊢ int¹ code)
let w = <if y == 0 then "hello" else x>
```

This let binding declares a normal variable `w` which depends on variables `x` and `y`, and produces a value of type `int¹ code`. The following function `f` takes an argument of this type, substituting `x` with `"world"` and `y` with `42`:

```
val f : (x : string¹; y : int¹ ⊢ int¹ code) → int¹ code
let f z = (z with x = "world"; y = 42)
```

We can then apply `f` to `w`:

```
f w  -- <if 42 == 0 then "hello" else "world">
```

We refer to `w` as an *unhygienic value*, and `f` as an *unhygienic function*; we will revisit this in §2.4.

Moreover, we note that dependencies can also be nested, allowing dependencies themselves to depend on other dependencies. For example, we can define:

```
val z : (s : (x : int¹ ⊢ string¹) ⊢ string¹ code)
let z = <cat (s with x = 2) (s with x = 2 + 1)>
```

Here, `z` depends on `s`, which in turn has a dependency on `x`. When `z` is used, such as in `z with s = e`, `s` is replaced by `e`, with `e` potentially using `x`. This `x` is subsequently replaced by `2` and `2 + 1`, respectively. As a concrete example, we have:

```
(z with s = string_of_int x)  -- <cat (string_of_int 2) (string_of_int (2 + 1))>
```

Lastly, a type can express variable dependencies across multiple levels, and the corresponding term can take expressions from these levels, as in the following program:

```
val h : (x : bool¹; y : int² ⊢ string² code code)
let h = <if x then <string_of_int y> else <string_of_int (y + 1)>>
```

2.4 Unhygienic Functions

Hygienic macros [Kohlbecker et al. 1986] guarantee that macro expansion does not accidentally capture variables. While hygienic macros are well-established, they can sometimes also be insufficient. Barzilay et al. [2011] observed common kinds of practically useful, yet *unhygienic*, macro patterns, particularly those which can implicitly introduce bindings. Two common examples include a looping macro (e.g. `while`) that implicitly binds a variable (e.g. `abort`) that can be used to escape the loop inside the loop body [Clinger 1991], and *anaphoric conditionals* which introduce a binding

to hold the value of the tested expression. Barzilay et al. [2011] described these unhygienic macros as *"notoriously difficult to deal with"*.

In this paper, we use *unhygienic functions* to mean functions whose arguments may depend on additional, *later-stage* variables provided when the function is used. These variables are always refined by explicit substitutions, reflecting how a term can declare its free variables in its type. We show how our system supports unhygienic functions within this context. Considering anaphoric conditionals as an example, we would like to create a function `aif`, with which we can write the following program:

```
aif <big_long_calculation> <foo it ...> <bar it ...>
```

Here, both the then- and else-branches use the variable `it` to represent the result of `big_long_calculation`. The program will expand to:

```
<let it = big_long_calculation in
 if it then (foo it ...) else (bar it ...)>
```

In statically-typed languages, `it` will obviously stand for `True` in the then-branch and `False` in the else-branch. In dynamically-typed languages like Scheme, however, `it` in the then-branch is not necessarily `True`, making the function more practical.[2] While we focus on anaphoric conditionals for their simplicity, the underlying principles extend to more complex examples like looping.

Using variable dependencies, we define `aif` with the following type signature. Note that the second and third arguments have an additional dependency on the variable `it`:

```
val aif : bool¹ code
       → (it : bool¹ ⊢ 'a¹ code) → (it : bool¹ ⊢ 'a¹ code) → 'a¹ code
```

When `aif` is applied, its type signature tells the type checker to introduce `it` into the scope of the second and third arguments, allowing them to directly use it.

Now that we have the type signature, we can implement `aif` as follows:

```
let aif cond foo bar =
  let$ s1 : bool¹ = cond in
  let$ s2 : (it : bool¹ ⊢ 'a¹) = (foo with it) in
  let$ s3 : (it : bool¹ ⊢ 'a¹) = (bar with it) in
  <let it = s1 in
   if it then (s2 with it) else (s3 with it)>
```

The `aif` function takes three code arguments: `cond`, `foo`, and `bar`, with the latter two depending on the variable `it`. The `let$` unwraps these code arguments and binds them to splice variables `s1`, `s2`, and `s3`, respectively, with `it` explicitly provided. The quoted expression then constructs the resulting code. Notably, while code is generated for both branches, only the branch corresponding to the value of `it` will be evaluated at runtime. Supporting unhygienic functions shows the expressiveness of our calculus, enabling a broader range of meta-programming patterns, including those that intentionally manipulate lexical scoping in a type-safe manner.

[2] In Scheme, expressions such as `(if 2 foo bar)` are valid, in which case `it` will be bound to `2` within the then-branch. Moreover, with mutable states, `it` can be dynamically bound to different values throughout program execution.

2.5 Pattern Matching on Code

We have so far focused on *generative* meta-programming, where smaller code fragments are combined to construct larger ones, as illustrated by the `power` and `aif` examples. *Analytic* meta-programming, in contrast, can inspect or deconstruct code, enabling useful techniques like code rewriting (§2.6) for optimization. This is often realized through pattern matching on code in staging calculi [Jang et al. 2022; Parreaux et al. 2017; Stucki et al. 2021].

We extend our calculus with support for pattern matching on code, integrating it seamlessly through our mechanism of splice variables. As an example, we can write the following program that simply swaps two addends:

```
val swap : int¹ code ⟶ int¹ code
let swap n =  match$ n with | (x `+ y) ⟶ <y + x>
                            | _ ⟶ n
```

The `match$` construct performs pattern matching on code. Patterns describe the structure *within* the quotation, distinguishing between *pattern variables* (e.g. `x` and `y`) that match any code expression, and identifier literals like `` `+ `` and `` `* `` that match those exact identifiers. Intuitively, pattern variables behave similarly to splice variables, as they match inside quotations and may only be used within quotations. In this case, therefore, we have `x : int¹` and `y : int¹`. When matching succeeds, we return `<y + x>`. As a result, we have:

```
swap <1 + 2>    -- generates <2 + 1>
```

This demonstrates how our calculus naturally supports code pattern matching using a mechanism similar to splice variables. In fact, code pattern matching can be viewed as a generalization of let-splice bindings (§6), with additional rules for typing code patterns.

One key challenge in code pattern matching, however, is handling open code. Specifically, consider the following match expression:

```
match$ <fun x ⟶ x + 1> with | (fun x ⟶ y) ⟶ e
```

A natural result would be for `y` to match on `x + 1`. However, this results in `y` storing a code fragment that refers to the bound variable `x`. When `y` is used in `e`, `x` is no longer in the scope! To address this issue, Stucki et al. [2018] require writing the pattern as `fun x ⟶ $(y(` x))`, which binds `y` to a function `fun x ⟶ <$(x) + 1>` with type `int¹ code ⟶ int¹ code`. In other words, `y` has been eta-expanded to take the open variable (`x`) as an input. However, explicitly applying pattern variables to free variables is inconvenient, and binding the pattern variable to a function may be undesirable. Moreover, Stucki et al. [2018] restrict patterns to the simply-typed lambda calculus, as it remains unclear how the approach generalizes to quotes, splices, or matches themselves.[3]

In our calculus, just like pattern matching can be viewed as a generalization of let-splice bindings, pattern variables with open code work similarly to splice variables with variable dependencies. In the above example, our system binds `y` to `x + 1` with type `(x : int¹ ⊢ int¹)`. We can thus use `y` by providing specific `x`, e.g.

```
match$ <fun x ⟶ x + 1> with
  | (fun x ⟶ y) ⟶ <y with (x = 2)>  -- generates <2 + 1>
```

Here, crucially, `x` is not matched by name. Rather, the pattern matches any code expression with the function form, and we can refer to the bound variable as `x` within the function body. That is, pattern matching is *binding-aware*, and operates on the expression's binding structure.

[3]Contextual modal type systems [Jang et al. 2022] and similar calculi [Parreaux et al. 2017] inherently support open code. These systems behave quite differently, requiring the type of code to explicitly specify its open variable (see §9).

Pattern variables with open code offer a useful technique for code manipulation. As a larger example, consider a program that computes the partial derivative of an arithmetic expression as a code fragment. The following `partial` function recursively matches its argument `e`, generating code for its partial derivative with respect to an variable `var`:

```
val (+) (*) : int¹ → int¹ → int¹
val partial : (var : int¹ ⊢ int¹ code → int¹ code)
let rec partial e = match$ e with
    | (`var) → <1>
    | (g `+ h) →
      let$ dg = (partial with var) <g> in
      let$ dh = (partial with var) <h> in <dg + dh>
    | (g `* h) →
      let$ dg = (partial with var) <g> in
      let$ dh = (partial with var) <h> in <g * dh + h * dg>
    | _ → <0>
```

We can apply `partial` by providing `var` and the expression. For example, the following program:

```
let df : (x : int¹, y : int¹ ⊢ int¹ code) = (partial with var = x) <x * y + 1>
```

generates `<(1 * y + x * 0) + 0>`. We then use `df` by providing specific `x` and `y`, such as `df with x = 1, y = 2`, which produces `<(1 * 2 + 1 * 0) + 0>`.

Using pattern variables with open code, we can perform pattern matching under a binder. Specifically, consider computing the partial derivative of a let expression `let (y : int) = f in g`. Using the chain rule, the derivative is computed as:

$$\partial_x\, g(x, f(x)) = \partial_x\, g(x, y)\,|_{y=f(x)} + \partial_y\, g(x, y)\,|_{y=f(x)} \cdot \partial_x\, f(x)$$

This can be implemented with an additional pattern matching case:

```
match$ e with
  | ...
  | (let (y : int¹) = f in g) →
    let$ df = (partial with var) <f> in
    let$ dg1 : (y : int¹ ⊢ int¹) = (partial with var) <g with y> in
    let$ dg2 : (y : int¹ ⊢ int¹) = (partial with var = y) <g with y> in
    <let (y : int¹) = f in (dg1 with y) + (dg2 with y) * df>
```

Again, `y` is not matched by name, and the pattern matches any code expression with the let form. In this case, `g` is matched as a splice variable with a variable dependency on `y`. Within the pattern case, `dg1` computes the derivative of `g` with respect to the given variable `var`, `dg2` computes the derivative of `g` with respect to `y`, and `df` computes the derivative of `f`. The final expression combines these derivatives according to the chain rule.

2.6 Code Rewriting

Lastly, we support *code rewriting*, a primitive that replaces instances of a pattern within a target expression with a replacement expression. This is particularly useful for optimizing generated code, which often contains redundancies. For example, the `df` example using the `partial` function above generates `<(1 * y + x * 0) + 0>`, where the operations `1 *`, `* 0`, and `+ 0` can be simplified. Code rewriting is expressed as:

```
e_target rewrite p → e_replacement
```

where `p` is a code pattern, and `e_target` and `e_replacement` are code expressions. Similar to pattern matching, the pattern `p` can introduce pattern variables, potentially containing open code, which can be then used inside `e_replacement`.

As an example, the following program uses code rewriting to simplify the code generated by `df`:

```
let df_opt : (x : int¹, y : int¹ ⊢ int¹ code) =
  <(df with x; y)>
    rewrite (z `* `0) → <0>
    rewrite (z `+ `0) → <z>
    rewrite (`1 `* z) → <z>
```

Here, the first rewrite simplifies `z `* `0` to `<0>`, where `` `* `` and `` `0 `` match those exact identifiers, and `z` is a pattern variable that matches any code expression, which can be used within `<z>`. Thus, any code expression matching `z `* `0` is simplified to `0`. The other rewrites work similarly. The entire program simplifies the code `<(1 * y + x * 0) + 0>` to `<y>`.

As another example, we consider a program which finds all the local variables defined to be pairs, and rewrites any projections on these variables into direct component accesses [Parreaux et al. 2017]. The following program shows how to use `rewrite` to traverse a term `pgm` and perform these rewrites:

```
pgm rewrite (let p : int¹ * int¹ = (a, b) in body) →      -- a: int¹; b: int¹
  let$ body2 : (p : int¹ * int¹ ⊢ 't¹) =       -- body: (p: int¹ * int¹ ⊢ 't¹)
    <body with p>
      rewrite (`fst `p) → <a>
      rewrite (`snd `p) → <b>
  in <body2 with p = (a, b)>
```

The first `rewrite` identifies let bindings that define a variable `p` to be a pair of `a` and `b`. Here, `a` and `b` are pattern variables, and `body` is matched as a splice variable with a variable dependency on `p`. Within `body with p`, we rewrite every occurrence of `` `fst `p `` to `<a>`, and `` `snd `p `` to `<b>`, and bind its splice result to `body2`. The program then returns `<body2 with p = (a, b)>`, substituting `p` with `(a,b)` within `body2`. For example, given `<let y = (1+2, 3) in (fst y + snd y, y)>`, the body matches `<(fst y + snd y, y)>`. Then, `body2` binds the rewritten program `((1+2)+3, y)`, and then the substitution turns it into `<((1+2)+3, (1+2, 3))>`.

3 A Staging Calculus with Splice Variables

This section presents $\lambda^{\bigcirc\triangleright}$, a typed lambda calculus with staging, quotations, and let-splice bindings. We begin by introducing the syntax of types and expressions (§3.1-3.2), and then present the typing rules (§3.3). The dynamic semantics are introduced in the following section (§4).

3.1 Types and Typing Contexts

Types. Types are defined by the following grammar:

$$A, B \ ::= \ \mathbf{bool} \mid [\Delta \vdash A] \to B \mid \bigcirc A \mid \Delta \triangleright A$$

- **bool** represents booleans.
- $[\Delta \vdash A] \to B$ represents *unhygienic functions* from A to B, where the argument may additionally depend on variables in Δ. We write $A \to B$ as shorthand for $[\cdot \vdash A] \to B$.
- $\bigcirc A$ represents quoted expressions of type A, whose computations happen at the next stage.

- $\Delta \rhd A$ represents *unhygienic values* of type A with dependencies Δ. This type allows any expressions to depend on later-stage variables, enabling unhygienic values to be used as first-class citizens. For example, we can express function types like $(\Delta \rhd A) \to (\Delta' \rhd B)$.

We note that while $[\Delta \vdash A] \to B$ can be expressed as $(\Delta \rhd A) \to B$[4], we include the former in the language as it allows for more natural definitions of unhygienic functions, without requiring explicit wrapping and unwrapping (§3.2).

Contexts. $\lambda^{\bigcirc\rhd}$ uses typing contexts to track *variable dependencies*. These contexts enable unhygienic values and serve as the foundation for code pattern matching, which is introduced in §6. Contexts are defined by the following grammar:

$$\Gamma, \Delta \ ::= \ \cdot \mid \Gamma, x : [\Delta \vdash^n A]$$

Each variable x within a context is associated with three components:

- A *context* Δ, which tracks the variable dependencies of x. When Δ is empty, we write $x : A^n$ as shorthand for $x : [\cdot \vdash^n A]$.
- An integer *level* n at which x is bound. The level has the same meaning as in Davies [1996]: higher values correspond to later stages, while lower values correspond to earlier stages.
- A *type* A, which describes the type of x. As we will see, types themselves can contain further variable dependencies, thus allowing for nested variable dependencies.

The *restriction* of a context Γ to level n, written $\Gamma\restriction_n$, removes all variables in Γ with levels less than n.

Well-stagedness. We consider only *well-staged* types and contexts in our typing rules (§3.3).

$\boxed{A^n}$ *(Well-staged Types)*

$$\textsc{WS-Bool}\ \frac{}{\mathbf{bool}^n} \qquad \textsc{WS-CtxArr}\ \frac{\Delta^{n+1} \quad A^n \quad B^n}{([\Delta \vdash A] \to B)^n} \qquad \textsc{WS-Code}\ \frac{A^{n+1}}{(\bigcirc A)^n} \qquad \textsc{WS-Wrap}\ \frac{\Delta^{n+1} \quad A^n}{(\Delta \rhd A)^n}$$

$\boxed{\Gamma^n}$ *(Well-staged Contexts)*

$$\textsc{WS-Empty}\ \frac{}{(\cdot)^n} \qquad \textsc{WS-Nest}\ \frac{\Gamma^n \quad \Delta^m \quad A^m \quad m \geq n}{(\Gamma, x : [\Delta \vdash^m A])^n}$$

For type $[\Delta \vdash A] \to B$ and $\Delta \rhd A$ to be well-staged at level n, Δ must be well-staged at level $n+1$. A context Γ is well-staged at level n, if each entry $x : [\Delta \vdash^m A]$ has $m \geq n$. In other words, stage levels can only stay the same or increase as the nesting of [] becomes deeper. This allows us to express nested variable dependencies as well as variable dependencies across multiple levels (§2.3).

We often write A^n and Γ^n to indicate the levels of a type or a context. This notation binds more tightly than type constructors. For example, $[\Delta^{n+1} \vdash A^n] \to B^n$ is a well-staged type at level n. Importantly, context restriction preserves well-stagedness: if Γ^n, then $(\Gamma\restriction_m)^m$.

3.2 Expressions

We define the syntax of expressions in $\lambda^{\bigcirc\rhd}$ as follows.

$$\begin{array}{lll}
e ::= x_\sigma & \textit{(Variables)} & \sigma ::= \cdot \quad \textit{(Empty)} \\
\quad \mid \mathbf{true} \mid \mathbf{false} \mid \mathbf{if}\ e_1\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3 & \textit{(Booleans)} & \quad \mid \sigma, x \mapsto y \quad \textit{(Renaming)} \\
\quad \mid \lambda_\Delta x : A.\ e \mid e_1\ e_2 & \textit{(Functions)} & \quad \mid \sigma, x \mapsto e \quad \textit{(Substitution)} \\
\quad \mid \langle e \rangle \mid \mathbf{let}_\Delta \langle x : A \rangle = e_1\ \mathbf{in}\ e_2 & \textit{(Quote and Unquote)} & \\
\quad \mid \mathbf{wrap}_\Delta e \mid \mathbf{let}\,\mathbf{wrap}_\Delta\ x : A = e_1\ \mathbf{in}\ e_2 & \textit{(Wrap and Unwrap)} &
\end{array}$$

[4]We have $\mathit{wrapToArr} : ((\Delta \rhd A) \to B) \to ([\Delta \vdash A] \to B) \coloneqq \lambda f.\ \lambda_\Delta x : A.\ f\ (\mathbf{wrap}_\Delta x_{\mathrm{id}_\Delta})$; and $\mathit{arrToWrap} : ([\Delta \vdash A] \to B) \to ((\Delta \rhd A) \to B) \coloneqq \lambda f.\ \lambda x : (\Delta \rhd A).\ f\ (\mathbf{let}\,\mathbf{wrap}_\Delta\ y : A = x\ \mathbf{in}\ y_{\mathrm{id}_\Delta})$. Notably, $[\Delta \vdash A] \to B$ brings syntactic convenience, such as putting Δ into scope when typing the function argument (rule CtxApp in Fig. 1).

Within expressions, x_σ represents a variable x paired with a *delayed substitution* σ. This substitution maps the dependencies of x to expressions in the current context. The substitution is applied when x is substituted with a concrete expression (see §4.1 for the formal definition of substitution).

An unhygienic function $\lambda_\Delta x : A.\ e$ allows its argument x to depend on variables in Δ. An application $e_1\ e_2$ applies a function e_1 to an argument e_2. A quotation $\langle e \rangle$ quotes an expression e into a code expression. Dually, a let-splice binding $\mathbf{let}_\Delta \langle x : A \rangle = e_1\ \mathbf{in}\ e_2$ unquotes a code expression e_1 that can depend on variables in Δ, introducing a next-stage variable x with dependencies Δ, which can be used inside quotations within e_2. Similarly, $\mathbf{wrap}_\Delta e$ wraps an expression e with dependencies Δ, allowing e to refer to variables in Δ, while $\mathbf{let\,wrap}_\Delta\ x : A = e_1\ \mathbf{in}\ e_2$ unwraps a wrapped expression e_1 with dependencies Δ, binding it to a current-stage variable x in e_2.

Substitutions σ can contain two kinds of entries: $x \mapsto y$ renames a dependency x to another variable y, and $x \mapsto e$ maps a dependency x to an expression e. We write id_Γ for the identity substitution that maps each variable in Γ to the same variable, i.e. $x_1 \mapsto x_1, x_2 \mapsto x_2 \ldots$ for $x_i \in \Gamma$. Moreover, the notation $\sigma\restriction_n$ removes all entries in σ with levels less than n, similar to context restriction. Lastly, for well-typed substitutions (§3.3), we write $x^m_\Delta \mapsto y$ or $x^m_\Delta \mapsto e$ to indicate the dependencies and stage level of a substitution entry.

Mapping between concrete and abstract syntax. As examples, (1) `x with y = e1; z = e2` is expressed as $x_{y \mapsto e_1, z \mapsto e_2}$; (2) `fun x : (Δ ⊢ A) → e` as $\lambda_\Delta x : A.\ e$; (3) `let x : (Δ ⊢ A) = e1 in e2` as $(\lambda_\Delta x : A.\ e_2)\ e_1$; (4) `let$ x : (Δ ⊢ A) = e1 in e2` as $\mathbf{let}_\Delta \langle x : A \rangle = e_1\ \mathbf{in}\ e_2$.

3.3 Typing Rules

The typing judgment $\Gamma \vdash^n e : A$ assigns a type A to an expression e under the context Γ, at stage level n. We require the context Γ to contain distinct variables, and to be well-staged at level n. The typing rules then ensure that A is also well-staged at level n. The judgment depends on another judgment, $\Gamma \vdash \sigma : \Gamma'$, which types a substitution σ between contexts Γ and Γ'. We similarly require both Γ and Γ' to contain distinct variables and to be well-staged. Intuitively, the substitution maps variables in Γ' to variables or expressions typed under Γ. Thus, we refer to Γ' the *domain* of σ.

The typing rules are defined in Fig. 1.

Rule VarSubst type-checks variables. A variable $x : [\Delta \vdash^n A]$ can only be used at level n, and must be provided a substitution σ that maps each dependency variable in Δ to a variable or an expression with the corresponding type under Γ. If Δ is empty, as is the case for normal variables, then σ is also empty. The next rules True, False, and If are standard.

Rule CtxAbs types an unhygienic function, which allows its argument to depend on Δ. Notably, a dependency context Δ is always staged one level higher than its surrounding context. The rule puts $(x : [\Delta^{n+1} \vdash^n A])$ into the context to type-check the body. Rule CtxApp type-checks an application. If e_1 has type $[\Delta^{n+1} \vdash A] \to B$, the rule extends the context with Δ to type-check the argument e_2.

Rule Quote increases the level to $n+1$ and restricts the context to $\Gamma\restriction_{n+1}$ (§3.1) to type-check the quoted expression e. If e has type A, then $\langle e \rangle$ has the code type $\bigcirc A$. Rule LetQuote unquotes a code expression e_1, and binds it to a variable x at the next level. Note that e_1 is typed under the current context extended with Δ, and thus x represents an open code fragment that may additionally depend on variables in a context Δ. The rule put x in the context to type-check e_2.

To type-check $\Delta \triangleright A$, rule Wrap extends the context with Δ to type-check e. Note that unlike rule Quote, rule Wrap does not change the stage level of the expression. Intuitively, quote corresponds to the next operator [Davies 1996, 2017], while $\Delta \triangleright A$ can be seen as Δ contains assumptions about future times. If each assumption in Δ is true at their corresponding time, then A is true at the current time. Rule LetWrap unwraps a wrapped expression and binds it to a contextual variable $x : [\Delta \vdash^n A]$ when typing e_2.

$\boxed{\Gamma \vdash^n e : A}$ *(Expression Typing)*

$$\textsc{VarSubst}\ \frac{\Gamma \ni x : [\Delta \vdash^n A] \qquad \Gamma \vdash \sigma : \Delta}{\Gamma \vdash^n x_\sigma : A} \qquad \textsc{True}\ \frac{}{\Gamma \vdash^n \mathbf{true} : \mathbf{bool}} \qquad \textsc{False}\ \frac{}{\Gamma \vdash^n \mathbf{false} : \mathbf{bool}}$$

$$\textsc{If}\ \frac{\Gamma \vdash^n e_1 : \mathbf{bool} \qquad \Gamma \vdash^n e_2 : A \qquad \Gamma \vdash^n e_3 : A}{\Gamma \vdash^n \mathbf{if}\ e_1\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3 : A} \qquad \textsc{CtxAbs}\ \frac{\Gamma, x : [\Delta^{n+1} \vdash^n A] \vdash^n e : B}{\Gamma \vdash^n \lambda_\Delta x : A.\ e : [\Delta \vdash A] \to B}$$

$$\textsc{CtxApp}\ \frac{\Gamma \vdash^n e_1 : [\Delta^{n+1} \vdash A] \to B \qquad \Gamma, \Delta \vdash^n e_2 : A}{\Gamma \vdash^n e_1\ e_2 : B} \qquad \textsc{Quote}\ \frac{\Gamma\restriction_{n+1} \vdash^{n+1} e : A}{\Gamma \vdash^n \langle e \rangle : \bigcirc A}$$

$$\textsc{LetQuote}\ \frac{\Gamma, \Delta^{n+1} \vdash^n e_1 : \bigcirc A \qquad \Gamma, x : [\Delta \vdash^{n+1} A] \vdash^n e_2 : B}{\Gamma \vdash^n \mathbf{let}_\Delta \langle x : A \rangle = e_1\ \mathbf{in}\ e_2 : B} \qquad \textsc{Wrap}\ \frac{\Gamma, \Delta^{n+1} \vdash^n e : A}{\Gamma \vdash^n \mathbf{wrap}_\Delta e : \Delta \rhd A}$$

$$\textsc{LetWrap}\ \frac{\Gamma \vdash^n e_1 : \Delta \rhd A \qquad \Gamma, x : [\Delta \vdash^n A] \vdash^n e_2 : B}{\Gamma \vdash^n \mathbf{let\ wrap}_\Delta\ x : A = e_1\ \mathbf{in}\ e_2 : B}$$

$\boxed{\Gamma \vdash \sigma : \Gamma'}$ *(Substitution Typing)*

$$\textsc{S-Empty}\ \frac{}{\Gamma \vdash \cdot : \cdot} \qquad \textsc{S-Rename}\ \frac{\Gamma \vdash \sigma : \Gamma' \qquad \Gamma \ni y : [\Delta \vdash^m A]}{\Gamma \vdash (\sigma, x \mapsto y) : \Gamma', x : [\Delta \vdash^m A]} \qquad \textsc{S-Subst}\ \frac{\Gamma \vdash \sigma : \Gamma' \qquad \Gamma\restriction_m, \Delta \vdash^m e : A}{\Gamma \vdash (\sigma, x \mapsto e) : \Gamma', x : [\Delta \vdash^m A]}$$

Fig. 1. $\lambda^{\bigcirc\rhd}$ typing rules

For typing substitutions, rule S-Rename checks that renaming preserves the stage level and dependencies of a variable. Rule S-Subst checks that substitution maps an m-level variable x to an m-level expression e, where Γ is restricted to level m and then extended with Δ to type-check e. Given the typing rules, we have $\Gamma \vdash \mathrm{id}_\Gamma : \Gamma$ (§3.2). Also, if $\Gamma_2 \vdash \sigma : \Gamma_1$, then $\Gamma_2\restriction_n \vdash \sigma\restriction_n : \Gamma_1\restriction_n$.

4 Dynamics

This section presents a small-step, call-by-value operational semantics for $\lambda^{\bigcirc\rhd}$, based on substitution.

4.1 Substitution

We define a substitution function $e[\sigma]$, which applies a typed substitution σ to a typed expression e, and $\sigma_1[\sigma]$, which applies σ to another typed substitution σ_1. Since expressions and substitutions can contain each other, these functions are defined mutually. These functions perform *simultaneous substitution*, meaning that all open variables in the expression are replaced at once. As a consequence, the domain of σ must match the context of the expression being substituted (see also Theorem 4.2).

The definitions are presented below, with $\Gamma_1 \vdash^n e : A$ and $\Gamma_2 \vdash \sigma : \Gamma_1$ (Theorem 4.2). For space reasons, we show a selection of rules; the complete set of rules can be found in §A.

$\boxed{e[\sigma]}$ *(Expression Substitution)*

$$
\begin{aligned}
(x_{\sigma_1})[\sigma] &\coloneqq \begin{cases} y_{(\sigma_1[\sigma])} & \text{if } \sigma(x) = y, \\ e[\mathsf{id}_{\Gamma_2\restriction_m}, \sigma_1[\sigma]] & \text{if } \sigma(x) = e. \end{cases} \\
(\langle e \rangle)[\sigma] &\coloneqq \langle e[\sigma\restriction] \rangle \\
(\mathbf{let}_\Delta \langle x : A \rangle = e_1 \mathbf{\ in\ } e_2)[\sigma] &\coloneqq \mathbf{let}_\Delta \langle x : A \rangle = (e_1[\sigma, \mathsf{id}_\Delta]) \mathbf{\ in\ } (e_2[\sigma, x \mapsto x]) \\
(\mathbf{wrap}_\Delta e)[\sigma] &\coloneqq \mathbf{wrap}_\Delta (e[\sigma, \mathsf{id}_\Delta]) \\
(\mathbf{let\ wrap}_\Delta\ x : A = e_1 \mathbf{\ in\ } e_2)[\sigma] &\coloneqq \mathbf{let\ wrap}_\Delta\ x : A = (e_1[\sigma]) \mathbf{\ in\ } (e_2[\sigma, x \mapsto x])
\end{aligned}
$$

$\boxed{\sigma_1[\sigma]}$ *(Substitution on Substitutions)*

$$
\begin{aligned}
(\cdot)[\sigma] &\coloneqq \cdot & (\sigma_1, x \mapsto y)[\sigma] &\coloneqq \sigma_1[\sigma], x \mapsto \sigma(y) \\
& & (\sigma_1, x^m_\Delta \mapsto e)[\sigma] &\coloneqq \sigma_1[\sigma], x \mapsto e[\sigma\restriction_m, \mathsf{id}_\Delta]
\end{aligned}
$$

For expression substitutions $e[\sigma]$, most cases simply extend or restrict the substitution to match the contexts of the subexpressions and then recurse. The most interesting case is substitution on a variable, $(x_{\sigma_1})[\sigma]$. Intuitively, applying σ to both x and σ_1 should yield $\sigma(x)_{(\sigma_1[\sigma])}$. Thus, when $\sigma(x) = y$, the function simply returns $y_{(\sigma_1[\sigma])}$. When $\sigma(x) = e$, however, since only variables can have delayed substitutions, the function must further apply $\sigma_1[\sigma]$ to e. In other words, the delayed substitution is *applied* at this point. The termination of this process is less obvious and requires a subtle argument, which is discussed in §4.1.1. The extended identity substitutions are for proofs (Theorem 4.2); a practical implementation may omit them.

To substitute a substitution, $\sigma_1[\sigma]$ simply applies σ to each entry of σ_1. If an entry is a variable y, it is replaced with the corresponding entry $\sigma(y)$. Otherwise, if the entry is an expression e, the function applies σ to e, restricting and extending σ to match the level and context of e.

4.1.1 Termination. The substitution $e[\sigma]$ defined above is not structurally recursive on e, so its termination is not immediately obvious. The problematic case is the second clause of $(x_{\sigma_1})[\sigma]$:

$$(x_{\sigma_1})[\sigma] = e[\mathsf{id}_{\Gamma_2\restriction_m}, \sigma_1[\sigma]] \quad \text{if } \sigma(x) = e.$$

Here, the expression e is not a sub-expression of x_{σ_1}, but rather an element of the substitution σ. Therefore, we cannot argue for termination based solely on the size of the input expression. A similar issue arises in Boespflug and Pientka [2011], where termination of single substitution is shown by demonstrating that the size of the input expression's context decreases.

In our case, since we work with simultaneous substitution, all expression entries within the substitution must be considered. To prove termination, we define the *depth* of a substitution σ as the maximum context depth among all its expression entries:

$\boxed{\mathsf{depth}(\Gamma)}$ *(Context Depth)* $\boxed{\mathsf{depth}(\sigma)}$ *(Substitution Depth)*

$$
\begin{aligned}
\mathsf{depth}(\cdot) &\coloneqq 0 & \mathsf{depth}(\cdot) &\coloneqq 0 \\
\mathsf{depth}(\Gamma, x : [\Delta \vdash^m A]) &\coloneqq & \mathsf{depth}(\sigma, x \mapsto y) &\coloneqq \mathsf{depth}(\sigma) \\
\mathsf{depth}(\Gamma) \sqcup (\mathsf{depth}(\Delta) + 1) & & \mathsf{depth}(\sigma, x^m_\Delta \mapsto e) &\coloneqq \mathsf{depth}(\sigma) \sqcup (\mathsf{depth}(\Delta) + 1)
\end{aligned}
$$

Intuitively, the depth of a substitution bounds the number of times the problematic case can occur in a recursive call. To prove this, we show that the depth decreases in the problematic case and does not increase in other cases. Non-increase follows because extending with variable entries does not change the depth, as they are not counted, and restriction does not increase the depth since the depth is defined as a maximum. In the problematic case, we need to show that

$$\mathsf{depth}(\mathsf{id}_{\Gamma_2\restriction_m}, \sigma_1[\sigma]) < \mathsf{depth}(\sigma).$$

$\boxed{\Gamma \vdash e \longrightarrow e'}$ *(Call-by-value Reduction)*

$$\textsc{CtxAppAbs}\ \frac{}{\Gamma \vdash (\lambda_\Delta x : A.\ e_1)\ v_2 \longrightarrow e_1[\mathrm{id}_\Gamma, x \mapsto v_2]} \qquad \textsc{LetQuoteQuote}\ \frac{}{\Gamma \vdash \mathbf{let}_\Delta \langle x : A \rangle = \langle e_1 \rangle\ \mathbf{in}\ e_2 \longrightarrow e_2[\mathrm{id}_\Gamma, x \mapsto e_1]}$$

$$\textsc{IfTrue}\ \frac{}{\Gamma \vdash \mathbf{if}\ \mathbf{true}\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3 \longrightarrow e_2} \qquad \textsc{IfFalse}\ \frac{}{\Gamma \vdash \mathbf{if}\ \mathbf{false}\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3 \longrightarrow e_3}$$

$$\textsc{LetWrapWrap}\ \frac{}{\Gamma \vdash \mathbf{let}\ \mathbf{wrap}_\Delta\ x : A = \mathbf{wrap}_\Delta\ v_1\ \mathbf{in}\ e_2 \longrightarrow e_2[\mathrm{id}_\Gamma, x \mapsto v_1]} \qquad \textsc{CongWrap}\ \frac{\Gamma, \Delta \vdash e \longrightarrow e'}{\Gamma \vdash \mathbf{wrap}_\Delta e \longrightarrow \mathbf{wrap}_\Delta e'}$$

Fig. 2. $\lambda^{\bigcirc\triangleright}$ operational semantics (excerpt)

Let Δ be the domain of σ_1. By induction, $\sigma_1[\sigma]$ is well-defined and $\Gamma_2 \vdash \sigma_1[\sigma] : \Delta$. Since adding variable entries does not change the depth, and since the depth of a substitution is at most the depth of its domain, we have $\mathrm{depth}(\mathrm{id}_{\Gamma_2\restriction_m}, \sigma_1[\sigma]) \le \mathrm{depth}(\sigma_1[\sigma]) \le \mathrm{depth}(\Delta)$. Since $x^m_\Delta \mapsto e \in \sigma$, it follows that $\mathrm{depth}(\Delta) < \mathrm{depth}(\sigma)$. Thus, we get $\mathrm{depth}(\mathrm{id}_{\Gamma_2\restriction_m}, \sigma_1[\sigma]) \le \mathrm{depth}(\Delta) < \mathrm{depth}(\sigma)$ as required. Therefore, the substitution function terminates and is well-defined on all inputs.

4.1.2 Substitution Properties. We prove that the substitution functions enjoy the usual associativity and commutativity properties. These lemmas are crucial when proving adequacy of the denotational semantics in §7.4, which employs a logical relations argument that famously relies on these properties [Abel et al. 2019; Acevedo and Weirich 2023].

Lemma 4.1 (Substitution Properties).

(Identity) *(1) If* $\Gamma_1 \vdash^n e : A$ *then* $e[\mathrm{id}_{\Gamma_1}] = e$.
(2) If $\Gamma_1 \vdash \sigma_1 : \Delta$ *then* $\sigma_1[\mathrm{id}_{\Gamma_1}] = \sigma_1 = \mathrm{id}_\Delta[\sigma_1]$.

(Associativity) *(1) If* $\Gamma_1 \vdash^n e : A$ *and* $\Gamma_2 \vdash \sigma : \Gamma_1$ *and* $\Gamma_3 \vdash \sigma' : \Gamma_2$ *then* $e[\sigma][\sigma'] = e[\sigma[\sigma']]$.
(2) If $\Gamma_1 \vdash \sigma_1 : \Delta$ *and* $\Gamma_2 \vdash \sigma : \Gamma_1$ *and* $\Gamma_3 \vdash \sigma' : \Gamma_2$ *then* $\sigma_1[\sigma][\sigma'] = \sigma_1[\sigma[\sigma']]$.

(Commutativity) *If* $\Gamma, \Delta \vdash^n e : A$, $\Gamma' \vdash \sigma_1 : \Gamma$, *and* $\Gamma' \vdash \sigma_2 : \Delta$, *then* $e[\sigma_1, \sigma_2] = e[\sigma_1, \mathrm{id}_\Delta][\mathrm{id}_{\Gamma'}, \sigma_2]$.

Since substitution functions are not structurally recursive, some of these properties need to be proven by induction on depths. Associativity, for instance, is proven by induction on sum of the depths of both substitutions, i.e. $\mathrm{depth}(\sigma) + \mathrm{depth}(\sigma')$. The full proofs are available in Agda (§8).

Lastly, we show that substitution preserves typing:

Lemma 4.2 (Substitution Lemma).

- *If* $\Gamma_1 \vdash^n e : A$ *and* $\Gamma_2 \vdash \sigma : \Gamma_1$ *then* $\Gamma_2 \vdash^n e[\sigma] : A$.
- *if* $\Gamma_1 \vdash \sigma_1 : \Delta$ *and* $\Gamma_2 \vdash \sigma : \Gamma_1$ *then* $\Gamma_2 \vdash \sigma_1[\sigma] : \Delta$.

4.2 Operational Semantics

We now present our reduction rules. First, values and evaluation contexts are defined below.

$$\begin{array}{lrcl} \text{Values} & v & ::= & \mathbf{true} \mid \mathbf{false} \mid \lambda_\Delta x : A.\ e \mid \langle e \rangle \mid \mathbf{wrap}_\Delta\ v \\ \text{Evaluation Contexts} & E & ::= & [\,] \mid E\ e_2 \mid v_1\ E \mid \mathbf{if}\ E\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3 \mid \mathbf{let}_\Delta \langle x : A \rangle = E\ \mathbf{in}\ e_2 \\ & & & \mid \mathbf{wrap}_\Delta E \mid \mathbf{let}\ \mathbf{wrap}_\Delta\ x : A = E\ \mathbf{in}\ e_2 \end{array}$$

Evaluation contexts are expressions with a hole ([]) that can be filled with another expression, indicating the part of the expression being reduced. Note that unlike Calcagno et al. [2003]; Taha and Sheard [2000], our definition of values does not require level indexing.

The call-by-value reduction is defined in Fig. 2. We write $\Gamma \vdash e \longrightarrow e'$ for a reduction step from e to e' under context Γ. The context is needed when applying substitutions (e.g. in rule LetQuoteQuote). Rule CtxAppAbs applies a lambda to an argument value by substituting the argument into the lambda body. Rule LetQuoteQuote unwraps a quoted expression $\langle e_1 \rangle$ and substitutes it into the body e_2; since any quoted expression is a value, e_1 itself is not necessarily a value. Rules IfTrue and IfFalse reduce an if-then-else expression to the appropriate branch. Rule LetWrapWrap unwraps a wrapped value $\mathbf{wrap}_\Delta\, v$ and substitutes it into the body e_2. Unlike rule LetQuoteQuote, the wrapped value is first reduced to a value before substitution. For each case of the evaluation context except for $[]$, there is a congruence rule that reduces the expression contained within it. For example, rule CongWrap reduces the subexpression within a wrap expression. The complete set of reduction rules can be found in §A.3.

Type Soundness. We prove preservation, which is a corollary of the substitution lemma (4.2):

Theorem 4.3 (Preservation). *If* $\Gamma \vdash^n e : A$ *and* $\Gamma \vdash e \longrightarrow e'$ *then* $\Gamma \vdash^n e' : A$.

Progress is subtler. Specifically, since we allow arbitrary nesting of dependencies, delayed substitutions are crucial for progress to hold in our calculus. For example, consider the program:

$$\mathbf{let}_{x:[z:\mathbf{bool} \vdash^1 \mathbf{bool}]} \langle y : \mathbf{bool} \rangle = (\mathbf{let}\, \langle z : \mathbf{bool} \rangle = \langle \mathbf{true} \rangle\, \mathbf{in}\, \langle x_{z \mapsto z} \rangle)\, \mathbf{in}\, \langle \mathbf{true} \rangle$$

Here, y is declared with a dependency x, which in turn depends on z. To evaluate the inner let binding, we need to substitute z with **true** within the let body. Without allowing delayed substitutions to contain arbitrary (open) expressions (e.g. $x_{z \mapsto \mathbf{true}}$), this substitution would do nothing, and then the evaluation would be stuck. In contrast, the core calculus of Xie et al. [2022] does not allow nested dependencies. As a result, in their system, variables can only capture their variable dependencies from the context.

We prove that progress holds for expressions that do not contain variables at the current level. This reflects our definition of "unhygienic values": values in this calculus are not necessarily closed terms, since quotations may include variables from later stages.

Theorem 4.4 (Progress). *If* $\Gamma^{n+1} \vdash^n e : A$ *then either* e *is a value or there exists* e' *such that* $\Gamma^{n+1} \vdash e \longrightarrow e'$.

4.3 Example

We demonstrate the operational semantics of the calculus with a larger example. For improved readability, we present the example using the concrete syntax; recall the mapping between the concrete syntax and the abstract syntax provided in §3.2. Consider the following program:

```
let$ y : (x : (z : bool¹ ⊢ bool¹) ⊢ bool¹) =                    1
  let$ z = <true> in <x with z = z>                             2
in                                                              3
  let$ x : (z : bool¹ ⊢ bool¹) = <not z> in                     4
  let$ z = <false> in                                           5
  <(y with x = x) and z>                                        6
```

The program first defines `y`, using the same definition as in §4.2. In line 4-5, it defines `x`, which depends on `z`, and then defines `z`. Finally, the program returns the code of the `and` operator applied to `y with x = x` and `z`. In line 6, which definition of `z` is supplied to `x`?

To answer this, let us go through the reduction steps. First, line 2 reduces to `<x with z = true>`, as discussed in the previous subsection.

```
let$ y : (x : (z : bool¹ ⊢ bool¹) ⊢ bool¹) =
  <x with z = true>
in ...
```

Then, `y` is substituted with `<x with z = true>` in the let body, where the delayed substitution `with x = x` is applied.

```
let$ x : (z : bool¹ ⊢ bool¹) = <not z> in
let$ z = <false> in
<(x with z = true) and z>
```

Next, `x` is substituted with `<not z>`. The delayed substitution `with z = true` is then applied, yielding `<not true>`.

```
let$ z = <false> in
<(not true) and z>
```

Finally, `z` is substituted with `<false>`, resulting in the value:

```
<(not true) and false>
```

Therefore, the answer to our earlier question is that `x` uses `z = true`.

What if we wanted `x` to use the definition `z = false` instead? We can achieve that by either modifying the definition of `y` to explicitly capture `z`:

```
let$ y : (x : (z : bool¹ ⊢ bool¹); z : bool¹ ⊢ bool¹) =
  <x with z = z>
in ...
  <(y with x = x; z = z) and z>
```

or by changing the definition of `y` to capture a non-capturing version of `x`:

```
let$ y : (x : bool¹ ⊢ bool¹) =
  <x>
in ...
  <(y with x = (x with z = z)) and z>
```

These examples illustrate the ability of our type system to express and enforce different kinds of variable dependencies within unhygienic programs.

5 Translation Between $\lambda^{\bigcirc\triangleright}$ and Linear-Time Temporal Logic

This section presents translations between $\lambda^{\bigcirc\triangleright}$, Davies [1996]'s staging calculus $\lambda^{\bigcirc}$, and Kojima and Igarashi [2011]'s constructive linear-time temporal logic (CLTL). §5.1 presents a type-preserving translation from $\lambda^{\bigcirc\triangleright}$ to $\lambda^{\bigcirc}$, by converting variable dependencies into function arguments. Building on this, §5.2 shows soundness of $\lambda^{\bigcirc\triangleright}$ with respect to CLTL. Then, §5.3 establishes completeness of $\lambda^{\bigcirc\triangleright}$ with respect to CLTL, by proving CLTL axioms as $\lambda^{\bigcirc\triangleright}$ terms. Thus, $\lambda^{\bigcirc\triangleright}$ is sound and complete with respect to CLTL. Finally, §5.4 briefly discusses a direct translation from $\lambda^{\bigcirc}$ to $\lambda^{\bigcirc\triangleright}$ following the approach in Xie et al. [2022].

5.1 Translation from $\lambda^{\bigcirc\triangleright}$ to $\lambda^{\bigcirc}$

Type Translation. We first translate $\lambda^{\bigcirc\triangleright}$ types into $\lambda^{\bigcirc}$ types, which involves converting variable dependencies into function arguments. Since the input and output of a function are staged at the same level, we add the appropriate number of $\bigcirc$ to lower the input type to the same level as the

output. For example, $(x : A^1, y : B^2) \triangleright C^0$ becomes $\bigcirc A \to \bigcirc\bigcirc B \to C$. The formal translation is defined as follows:

$\boxed{[\![A^n]\!]}$ *(Type Translation)*

$$[\![([\Delta \vdash A] \to B)^n]\!] \coloneqq (\Delta \downarrow^n [\![A^n]\!]) \to [\![B^n]\!] \qquad [\![\mathbf{bool}^n]\!] \coloneqq \mathbf{bool}$$
$$[\![(\Delta \triangleright A)^n]\!] \coloneqq \Delta \downarrow^n [\![A^n]\!] \qquad [\![(\bigcirc A)^n]\!] \coloneqq \bigcirc [\![A^{n+1}]\!]$$

$\boxed{\Gamma \downarrow^n A}$ *(Dependency Translation)*

$$\cdot \downarrow^n A \coloneqq A \qquad (\Gamma, x : [\Delta \vdash^m A]) \downarrow^n B \coloneqq \Gamma \downarrow^n \left(\bigcirc^{m-n}(\Delta \downarrow^m [\![A]\!]) \to B\right)$$

The notation $(\Gamma \downarrow^n A)$ converts the dependencies in Γ into function arguments of a function with output type A, where n is the level of the output type. Since Γ can be nested, each entry in Γ is recursively translated at its respective level m, then lowered to level n by wrapping it in $\bigcirc^{m-n}$. Specifically, if $\Gamma = x_1 : [\Delta_1 \vdash^{m_1} A_1], \ldots, x_k : [\Delta_k \vdash^{m_k} A_k]$, then $\Gamma \downarrow^n A = \bigcirc^{m_1-n}(\Delta_1 \downarrow^{m_1} [\![A_1]\!]) \to \cdots \to \bigcirc^{m_k-n}(\Delta_k \downarrow^{m_k} [\![A_k]\!]) \to A$.

Context Translation. The translation from $\lambda^{\bigcirc\triangleright}$ contexts to $\lambda^{\bigcirc}$ contexts is given below, where dependencies of each variable are translated using the dependency translation.

$\boxed{[\![\Gamma]\!]}$ *(Context Translation)*

$$[\![\cdot]\!] \coloneqq \cdot \qquad [\![\Gamma, x : [\Delta \vdash^m A]]\!] \coloneqq [\![\Gamma]\!], x : (\Delta \downarrow^m [\![A]\!])^m$$

Term Translation. We now define a term translation from $\lambda^{\bigcirc\triangleright}$ to $\lambda^{\bigcirc}$. The translation preserves types but may introduce additional beta redexes inside quotations, similar to the example in §2.2. We present the translation with selected cases below, where $\langle e \rangle^n$ quotes e by n times, $\$^n(e)$ splices e by n times, $(\boldsymbol{\lambda}\Delta.\ e)$ abstracts an unhygienic term e with respect to Δ using lambda abstractions, and $(x \bullet \sigma)$ applies a variable x to each translated element in σ.

$\boxed{[\![e]\!]}$ *(Expression Translation)*

$$[\![x_\sigma]\!] \coloneqq x \bullet \sigma \qquad [\![\mathbf{let}_\Delta \langle x : A \rangle = e_1 \mathbf{\ in\ } e_2]\!] \coloneqq \mathbf{let}\ x = \langle \boldsymbol{\lambda}\Delta.\ \$[\![e_1]\!] \rangle \mathbf{\ in\ } ([\![e_2]\!][\$(x)/x])$$
$$[\![\lambda_\Delta x : A.\ e]\!] \coloneqq \lambda x.\ [\![e]\!] \qquad [\![\mathbf{wrap}_\Delta e]\!] \coloneqq \boldsymbol{\lambda}\Delta.\ [\![e]\!]$$
$$[\![e_1\ e_2]\!] \coloneqq [\![e_1]\!]\ (\boldsymbol{\lambda}\Delta.\ [\![e_2]\!]) \qquad [\![\mathbf{let\ wrap}_\Delta\ x : A = e_1 \mathbf{\ in\ } e_2]\!] \coloneqq \mathbf{let}\ x = [\![e_1]\!] \mathbf{\ in\ } [\![e_2]\!]$$
$$[\![\langle e \rangle]\!] \coloneqq \langle [\![e]\!] \rangle$$

$\boxed{\boldsymbol{\lambda}\Delta.\ e}$ *(Dependency Abstraction)*

$$\boldsymbol{\lambda}(\cdot).\ e \coloneqq e$$
$$\boldsymbol{\lambda}(\Delta, x : [\Delta' \vdash^m A]).\ e \coloneqq \boldsymbol{\lambda}\Delta.\ (\lambda x.\ e[\$^{m-n}\, x/x])$$

$\boxed{x \bullet \sigma}$ *(Dependency Application)*

$$x \bullet (\cdot) \coloneqq x$$
$$x \bullet (\sigma, y^m_\Delta \mapsto e) \coloneqq (x \bullet \sigma)\ \langle \boldsymbol{\lambda}\Delta.\ [\![e]\!] \rangle^{m-n}$$
$$x \bullet (\sigma, y^m_\Delta \mapsto z) \coloneqq (x \bullet \sigma)\ \langle z \rangle^{m-n}$$

We prove that the translation preserves typing:

Lemma 5.1 ($[\![\cdot]\!]$ preserves typing). *If* $\Gamma \vdash^n e : A$ *in* $\lambda^{\bigcirc\triangleright}$ *then* $[\![\Gamma]\!] \vdash^n [\![e]\!] : [\![A]\!]$ *in* $\lambda^{\bigcirc}$.

5.2 Translation from $\lambda^{\bigcirc\triangleright}$ to CLTL

CLTL [Kojima and Igarashi 2011] is a Hilbert-style axiomatization of $\lambda^{\bigcirc}$, which is characterized by the following axioms and rules. Note that in Hilbert-style systems, inference rules apply only to judgments with empty contexts; we denote this by writing $\vdash A$ in the rules:

Axioms (1) any intuitionistic tautology instance
(2) $\mathbf{K} : \bigcirc(A \to B) \to \bigcirc A \to \bigcirc B$ (3) $\mathbf{CK} : (\bigcirc A \to \bigcirc B) \to \bigcirc(A \to B)$

Rules (1) $\dfrac{\vdash A \to B \qquad \vdash A}{\vdash B}$ *(Modus Ponens)* (2) $\dfrac{\vdash A}{\vdash \bigcirc A}$ *(Necessitation)*

Kojima and Igarashi proved that $\lambda^{\bigcirc}$ is sound and complete with respect to CLTL. Therefore, by Lemma 5.1, $\lambda^{\bigcirc\triangleright}$ is also sound with respect to CLTL. In particular, the following translation from $\lambda^{\bigcirc\triangleright}$ judgments to CLTL formulas holds:

Lemma 5.2 (Soundness w.r.t. CLTL). *If* $\Gamma \vdash^n e : A$ *for some* e *in* $\lambda^{\bigcirc\triangleright}$, *then* $\vdash \Gamma \Downarrow^n [\![A]\!]$ *in CLTL.*

The translation provides temporal interpretation for the $\lambda^{\bigcirc\triangleright}$ types: $\bigcirc$ corresponds to the *next* operator, while $\Delta \triangleright A$ asserts that A holds at the current time if each assumption in Δ holds at its respective time.

5.3 CLTL Axioms as $\lambda^{\bigcirc\triangleright}$ Terms

Next, we show completeness of $\lambda^{\bigcirc\triangleright}$ with respect to CLTL through a reverse translation. Specifically, the axioms of CLTL [Kojima and Igarashi 2011] can be proven by the following $\lambda^{\bigcirc\triangleright}$ terms:

$$\mathbf{K} : \bigcirc(A \to B) \to \bigcirc A \to \bigcirc B \coloneqq \lambda f.\ \lambda x.\ \mathbf{let}\ \langle f' : A \to B \rangle = f\ \mathbf{in}\ \mathbf{let}\ \langle x' : A \rangle = x\ \mathbf{in}\ \langle f'\ x' \rangle$$
$$\mathbf{CK} : (\bigcirc A \to \bigcirc B) \to \bigcirc(A \to B) \coloneqq \lambda f.\ \mathbf{let}_{x:A^{n+1}}\ \langle y : B \rangle = f\ \langle x \rangle\ \mathbf{in}\ \langle \lambda x.\ y_{x \mapsto x} \rangle$$

where A and B are types staged at level $n+1$. Notably, the ability to introduce dependencies in rule LetQuote is crucial for proving **CK**. Moreover, the rules of CLTL are special cases of rule CtxApp and rule Quote with empty contexts. We have thus demonstrated that $\lambda^{\bigcirc\triangleright}$ is complete with respect to CLTL.

Lemma 5.3 (Completeness w.r.t. CLTL). *If* $\vdash A$ *in CLTL, then* $\cdot \vdash^0 e : A$ *for some* e *in* $\lambda^{\bigcirc\triangleright}$.

With Lemma 5.2, $\lambda^{\bigcirc\triangleright}$ is sound and complete with respect to CLTL.

5.4 Translation from $\lambda^{\bigcirc}$ to $\lambda^{\bigcirc\triangleright}$

A direct translation from $\lambda^{\bigcirc}$ to $\lambda^{\bigcirc\triangleright}$ can be achieved using a translation similar to the one described by Xie et al. [2022] for translating $\lambda^{\bigcirc}$ to their $F^{[\![]\!]}$ calculus. Their translation lifts splice expressions and binds them as *splice environments* associated with the innermost enclosing quotation at the same level. Because $\lambda^{\bigcirc\triangleright}$ generalizes the splice environments of $F^{[\![]\!]}$, their translation can be readily adapted. In our case, this involves lifting the splice expressions and binding them as let-splice bindings instead. The lifting process captures all encountered lambda bindings as variable dependencies, and these splice variables are provided with delayed substitutions. As an example,

$$\langle \lambda y.\ \$(power\ \langle y \rangle\ 2) + 3 \rangle \quad \text{translates to} \quad \mathbf{let}_{(y:\mathbf{int})}\ \langle z : \mathbf{int} \rangle = power\ \langle y \rangle\ 2\ \mathbf{in}\ \langle \lambda y : \mathbf{int}.\ z_{y \mapsto y} + 3 \rangle$$

Here, the splice $\$(power\ \langle y \rangle\ 2)$ is lifted to become a splice variable z outside of the quotation. However, since it refers to a local variable y, the splice variable z includes $y : \mathbf{int}$ in its variable dependency. When z is subsequently used inside the quotation, we provide the substitution for y. Once again, the ability to introduce variable dependencies plays a crucial role in the translation. For a more detailed explanation, we refer the reader to Xie et al. [2022].

6 Code Pattern Matching

This section introduces $\lambda^{\bigcirc\triangleright}_{\text{pat}}$, an extension of $\lambda^{\bigcirc\triangleright}$ with code pattern matching and code rewriting. We present its syntax (§6.1), typing rules (§6.2), and operational semantics (§6.3).

6.1 Syntax

We extend the syntax of $\lambda^{\bigcirc\triangleright}$ with two new expression forms: if-let expressions for code pattern matching and rewrite expressions for code rewriting.

$$e ::= \ldots \mid \textbf{if let}_\Delta \langle p \rangle = e_1 \textbf{ then } e_2 \textbf{ else } e_3 \mid e_1 \textbf{ rewrite } \langle p \rangle \mapsto e_2 \qquad \Gamma, \Delta, \Pi ::= \cdot \mid \Gamma, x : [\Delta \vdash^n A]$$
$$p ::= \hat{x} : A \mid \textit{(inherits every production of } e\textit{)} \qquad \pi ::= \cdot \mid \pi, x \mapsto y \mid \pi, x \mapsto p$$

$\textbf{if let}_\Delta \langle p \rangle = e_1 \textbf{ then } e_2 \textbf{ else } e_3$ matches the code expression e_1 against the pattern p. This can be seen as a generalization of the $\textbf{let}_\Delta \langle x : A \rangle = e_1 \textbf{ in } e_2$ expression in $\lambda^{\bigcirc\triangleright}$, where $x : A$ becomes a general pattern p and thus the match always succeeds. Inside the if-let expression, if the match succeeds, e_2 is evaluated with the pattern variables in p bound to the match results. Otherwise, e_3 is evaluated, with the pattern variables unavailable. The multi-branch `match$` expression used in the concrete syntax can be desugared into nested if-let expressions. Compared to multi-branch matching, if-let expressions are more convenient for formalization and ensure matching totality. $e_1 \textbf{ rewrite } \langle p \rangle \mapsto e_2$ replaces occurrences of the pattern p within e_1 with e_2, where pattern variables in p match sub-expressions in e_1 and become available within e_2.

Code patterns p are expressions with pattern variables. As syntactical notations to distinguish pattern variables from regular variables, we use $\hat{x}$ to denote pattern variables and y for regular variables. For example, a pattern $\lambda y : A.\ \hat{x} + y_\sigma$ contains a pattern variable $\hat{x}$ and a binding y. For simplicity, we assume pattern variables are distinct within a pattern. Substitution patterns π are used to match substitutions, where each entry can be either a variable or a pattern. Lastly, we write Π to denote contexts of pattern variables, whose definitions are the same as regular contexts Γ and Δ.

6.2 Typing Rules

We extend expression typing $\Gamma \vdash^n e : A$ with the following two rules:

$\boxed{\Gamma \vdash^n e : A}$ *(Expression Typing (extended))*

$$\textsc{IfLet}\ \frac{\Gamma\restriction_{n+1}; \Delta^{n+1} \vdash^{n+1} p : A \rightsquigarrow \Pi^{n+1} \qquad \Gamma, \Delta \vdash^n e_1 : \bigcirc A \qquad \Gamma, \Pi \vdash^n e_2 : B \qquad \Gamma \vdash^n e_3 : B}{\Gamma \vdash^n \textbf{if let}_\Delta \langle p \rangle = e_1 \textbf{ then } e_2 \textbf{ else } e_3 : B}$$

$$\textsc{Rewrite}\ \frac{\Gamma \vdash^n e_1 : \bigcirc A \qquad \Gamma\restriction_{n+1}; \cdot \vdash^{n+1} p : B \rightsquigarrow \Pi^{n+1} \qquad \Gamma, \Pi \vdash^n e_2 : \bigcirc B}{\Gamma \vdash^n e_1 \textbf{ rewrite } \langle p \rangle \mapsto e_2 : \bigcirc A}$$

Rule IfLet first type-checks the pattern p, yielding a type A and a pattern variable context Π^{n+1} (the typing rules for patterns are explained below). It then checks that e_1 has type $\bigcirc A$, and extends the typing context with Π^{n+1} when checking e_2. The rule can be seen as a generalization of rule LetQuote, which extends the context with only a single variable ($x : [\Delta \vdash^{n+1} A]$) when checking e_2. Finally, the expressions e_2 and e_3 must have the same type.

Rule Rewrite is similar. It first type-checks the pattern p, and extends the typing context with Π^{n+1} when checking e_1. The rule ensures that both e_1 and e_2 have code types. While the expression e_2 can have any type $\bigcirc B$, only sub-expressions with the same type as p are considered for rewriting.

The pattern typing judgment $\Gamma; \Delta \vdash^n p : A \rightsquigarrow \Pi$ checks the pattern p under contexts Γ and Δ, producing a type A and a context Π of pattern variables. The judgment uses two typing contexts: the expression context Γ contains variables from the surrounding context of the expression (see rules IfLet and Rewrite), allowing patterns to refer to existing variables, while the local context Δ contains local variables introduced either by $\textbf{let}_\Delta$ (rule IfLet) or within the pattern p. Separating these contexts ensures that each pattern variable captures the correct dependencies. For example,

in $\langle(\lambda x.\ \hat{y})\ \hat{z}\rangle$, the pattern variable $\hat{y}$ will capture x since it matches on a sub-expression that may contain x, while $\hat{z}$ will not. We present selected typing rules below.

$\boxed{\Gamma; \Delta \vdash^n p : A \rightsquigarrow \Pi}$ *(Code Pattern Typing (excerpt))*

$$\textsc{P-PVar}\ \frac{}{\Gamma; \Delta \vdash^n (\hat{x} : A) : A \rightsquigarrow x : [\Delta \vdash^n A]} \qquad \textsc{P-VarSubst1}\ \frac{\Gamma \ni x : [\Delta' \vdash^n A] \qquad \Gamma, \Delta \vdash \sigma : \Delta'}{\Gamma; \Delta \vdash^n x_\sigma : A \rightsquigarrow \cdot}$$

$$\textsc{P-VarSubst2}\ \frac{\Delta \ni x : [\Delta' \vdash^n A] \qquad \Gamma; \Delta \vdash \pi : \Delta' \rightsquigarrow \Pi}{\Gamma; \Delta \vdash^n x_\pi : A \rightsquigarrow \Pi} \qquad \textsc{P-CtxAbs}\ \frac{\Gamma; \Delta, x : [\Delta'^{n+1} \vdash^n A] \vdash^n p : B \rightsquigarrow \Pi}{\Gamma; \Delta \vdash^n (\lambda_{\Delta'} x : A.\ p) : [\Delta' \vdash A] \to B \rightsquigarrow \Pi}$$

$$\textsc{P-CtxApp}\ \frac{\Gamma; \Delta \vdash^n p_1 : [\Delta' \vdash A] \to B \rightsquigarrow \Pi_1 \qquad \Gamma; \Delta, \Delta' \vdash^n p_2 : A \rightsquigarrow \Pi_2}{\Gamma; \Delta \vdash^n p_1\ p_2 : B \rightsquigarrow \Pi_1, \Pi_2}$$

Rule P-PVar types pattern variables, producing a single pattern variable that captures the local context Δ. Rule P-VarSubst1 matches variables in the expression context Γ using a constant substitution σ, thereby enabling matching from any expression context. This contrasts with contextual type systems [Jang et al. 2022; Parreaux et al. 2017], which restrict matching to variables from the context the code is closed at. Rule P-VarSubst2 matches variables in the local context Δ, allowing for further matching against the variable's associated substitution using a substitution pattern π. This distinction between rule P-VarSubst1 and rule P-VarSubst2 ensures distinct pattern variables under substitution. Specifically, variables in Γ may be substituted with arbitrary terms. For example, consider a pattern $\langle x_{y \mapsto \hat{z}} \rangle$ where $x \in \Gamma$ and $\hat{z}$ is a pattern variable. If x is substituted with a term where y is not used linearly, such as 0 or $y + y$, then $\hat{z}$ would appear zero or multiple times, respectively.[5]

In rule P-CtxAbs, the local variable x is added to Δ to check the pattern p. In rule P-CtxApp, the pattern variables produced by p_1 and p_2 are combined in the result.

Similarly, the judgment $\Gamma; \Delta \vdash \pi : \Gamma' \rightsquigarrow \Pi$ type-checks substitution patterns, producing a typing context Γ' and a context Π of pattern variables:

$\boxed{\Gamma; \Delta \vdash \pi : \Gamma' \rightsquigarrow \Pi}$ *(Substitution Pattern Typing)*

$$\textsc{P-S-Empty}\ \frac{}{\Gamma; \Delta \vdash \cdot : \cdot \rightsquigarrow \cdot} \qquad \textsc{P-S-Var}\ \frac{\Gamma; \Delta \vdash \pi : \Gamma' \rightsquigarrow \Pi \qquad \Gamma, \Delta \ni y : [\Delta' \vdash^m A]}{\Gamma; \Delta \vdash (\pi, x \mapsto y) : \Gamma', x : [\Delta' \vdash^m A] \rightsquigarrow \Pi}$$

$$\textsc{P-S-Pattern}\ \frac{\Gamma; \Delta \vdash \pi : \Gamma' \rightsquigarrow \Pi_1 \qquad \Gamma\restriction_m; \Delta\restriction_m, \Delta'^m \vdash^m p : A \rightsquigarrow \Pi_2}{\Gamma; \Delta \vdash (\pi, x \mapsto p) : \Gamma', x : [\Delta' \vdash^m A] \rightsquigarrow \Pi_1, \Pi_2}$$

Rule P-S-Empty is trivial. When the substitution entry maps to a variable y, rule P-S-Var ensures that y is in either Γ or Δ, and recursively types π. When the substitution entry maps to a pattern p, rule P-S-Pattern type-checks the pattern p, and combines Π_1 and Π_2.

[5] Allowing a pattern variable to appear multiple times within a pattern is straightforward, by simply requiring it to map to the same expression during the merging of Π contexts (e.g. in rule P-CtxApp). Conversely, a pattern variable must continue appearing at least once, as it can be used within the matching branch. Mœbius [Jang et al. 2022] declares all pattern variables upfront, and the operational semantics requires a substitution of every declared pattern variable when matching an expression. Consequently, their operational semantics is non-deterministic when a pattern variable is unused.

6.3 Operational Semantics

Substitutions and evaluation contexts are mostly straightforward extensions of those in $\lambda^{\bigcirc\triangleright}$.

$$E ::= \ldots \mid \textbf{if let}_\Delta \langle p \rangle = E \textbf{ then } e_2 \textbf{ else } e_3 \mid E \textbf{ rewrite } \langle p \rangle \mapsto e_2 \mid v_1 \textbf{ rewrite } \langle p \rangle \mapsto E$$

The reduction rules for if-let and rewrite expressions are defined as follows, which rely on the meta-level functions match (§6.3.1) and rewrite (§6.3.2), respectively.

$\boxed{\Gamma \vdash e_1 \longrightarrow e_2}$ *(Reduction (extended))*

IfLetQuote1
$$\frac{\mathsf{match}(p; e_1) = \sigma}{\Gamma \vdash \textbf{if let}_\Delta \langle p \rangle = \langle e_1 \rangle \textbf{ then } e_2 \textbf{ else } e_3 \longrightarrow e_2[\mathsf{id}_\Gamma, \sigma]}$$

IfLetQuote2
$$\frac{\mathsf{match}(p; e_1) \textbf{ undefined}}{\Gamma \vdash \textbf{if let}_\Delta \langle p \rangle = \langle e_1 \rangle \textbf{ then } e_2 \textbf{ else } e_3 \longrightarrow e_3}$$

RewriteQuoteQuote
$$\frac{}{\Gamma \vdash \langle e_1 \rangle \textbf{ rewrite } \langle p \rangle \mapsto \langle e_2 \rangle \longrightarrow \langle \mathsf{rewrite}_{\mathsf{BottomUp}}(e_1; p; e_2) \rangle}$$

We also define the corresponding congruence rules for if-let and rewrite expressions. The complete set of reduction rules is given in §B.3.

We have proven the preservation (Theorem 4.3) and progress (Theorem 4.4) theorems for $\lambda^{\bigcirc\triangleright}_{\text{pat}}$. The proofs are available in Agda (§8).

6.3.1 Pattern Matching. Matching is defined as partial functions match. We present a selection of the rules below. Note that $\mathsf{match}(p; e)$ is defined up to α-equivalence on e: that is, we allow renaming of bound variables in e. For contexts introduced by λ_Δ, $\textbf{let}_\Delta$, or $\textbf{if let}_\Delta$, renaming is allowed, but reordering is not. These align with the de Bruijn representation used in the formalization (§8).

$\boxed{\mathsf{match}(p; e), \mathsf{match}(\pi; \sigma)}$ *(Expression and Substitution Matching (excerpt))*

$$\mathsf{match}(\hat{x} : A; e) \coloneqq x \mapsto e \qquad \mathsf{match}((\lambda_\Delta x : A.\ p); (\lambda_\Delta x : A.\ e)) \coloneqq \mathsf{match}(p; e)$$

$$\mathsf{match}(x_\pi; x_\sigma) \coloneqq \mathsf{match}(\pi; \sigma) \qquad \mathsf{match}(p_1\ p_2; e_1\ e_2) \coloneqq \mathsf{match}(p_1; e_1), \mathsf{match}(p_2; e_2)$$

6.3.2 Rewriting. Rewriting builds on the matching function by applying it to sub-expressions within a target expression, replacing any matches with a specified replacement expression. Given a target expression $\Gamma \vdash^n e_1 : A$, a pattern $\Gamma; \cdot \vdash^n p : B \rightsquigarrow \Pi$, and a replacement expression $\Gamma, \Pi \vdash^n e_2 : B$, a meta-level function rewrite produces an expression of the same type as e_1.

The rewrite function can be implemented using various strategies. For example, $\mathsf{rewrite}_{\mathsf{BottomUp}}$ rewrites sub-expressions before rewriting their parent expression. We define a helper function $\mathsf{applyBottomUp}(F; e)$, which applies a meta-level function F to each sub-expression of e with same level as e, replacing them in bottom-up order. For instance, $\mathsf{applyBottomUp}(F; (\lambda x.\ f\ x)) = F(\lambda x.\ F(F(f)\ F(x)))$. The complete definition of applyBottomUp is given in §B.5.

$\boxed{\mathsf{rewrite}_{\mathsf{BottomUp}}(e_1; p; e_2)}$ *(Bottom-up Rewriting)*

$$\mathsf{rewrite}_{\mathsf{BottomUp}}(e_1; p; e_2) \coloneqq \mathsf{applyBottomUp}(F; e_1), \text{ where}$$

$$F(e) \coloneqq \begin{cases} e_2[\mathsf{id}_\Gamma, \sigma] & \text{if } p \text{ and } e \text{ have the same type and } \mathsf{match}(p; e) = \sigma, \\ e & \text{otherwise.} \end{cases}$$

At each sub-expression e of e_1, the function F checks if e matches the pattern p. If so, it substitutes e_2 with the match result σ and replace e with it. Otherwise, it leaves e unchanged. Other strategies, such as rewriting only the top-most occurrences of the pattern, can also be defined. We include the top-most rewriting strategy in the appendix and in our Agda development. We may also iterate the rewriting process until a fixpoint is reached [Parreaux et al. 2017].

Lastly, we note that in rule RewriteQuoteQuote, the replacement expression is evaluated into a value $\langle e_2 \rangle$, before $\text{rewrite}_{\text{BottomUp}}$ takes place. Thus, while the replacement expression may refer to pattern variables in p, its evaluation cannot access the specific values matched by p. In some cases, we may want to inspect the match results within the replacement expression (§2.6). This requires a different semantics where the replacement expression is evaluated *at each pattern match*, with the corresponding match results available. We defer the definition of this alternative semantics to §7.3, where it can be more naturally presented after the introduction of the denotational semantics.

7 Denotational Semantics

This section presents the denotational semantics for $\lambda^{\bigcirc\triangleright}$ and $\lambda^{\bigcirc\triangleright}_{\text{pat}}$, providing interpretations of types (§7.1), contexts (§7.2), and expressions (§7.3). We then define a logical relation (§7.4) and establish the fundamental property (Theorem 7.1), which implies termination of the operational semantics and the adequacy of the denotational semantics.

At a high-level, since evaluation happens at a specific stage n, we refer to this stage n as the *current stage*, and any stages from $n+1$ onward as *later stages*. Semantic interpretations are given to current-stage terms, while later-stage terms remain syntactic. We employ a Kripke-style model [Asai et al. 2014; Mitchell and Moggi 1991], where "worlds" are later-stage contexts Γ^{n+1} related by later-stage substitutions ${\Gamma'}^{n+1} \vdash \sigma : \Gamma^{n+1}$. Types and contexts are interpreted as sets indexed by later-stage contexts, where these sets are required to be *monotonic* with respect to the index. Formally, given an indexed set $D_{(\cdot)}$ and a later-stage substitution $\Gamma' \vdash \sigma : \Gamma$, there is a function $D_\Gamma \to D_{\Gamma'}$ mapping elements from D_Γ to $D_{\Gamma'}$.[6]

7.1 Type Interpretation

Interpretations of level-n types are defined by recursion on the type structure:

$\boxed{(\!|\, A^n \,|\!)_\Gamma}$ *(Type Interpretation)*

$$(\!|\, [\Delta \vdash A] \to B \,|\!)_\Gamma \coloneqq \textstyle\prod_{(\Gamma' \vdash \sigma : \Gamma)} (\!|\, A \,|\!)_{\Gamma', \Delta} \to (\!|\, B \,|\!)_{\Gamma'} \qquad (\!|\, \mathbf{bool} \,|\!)_\Gamma \coloneqq \{\mathsf{True}, \mathsf{False}\}$$

$$(\!|\, \Delta \triangleright A \,|\!)_\Gamma \coloneqq (\!|\, A \,|\!)_{\Gamma, \Delta} \qquad (\!|\, \bigcirc A \,|\!)_\Gamma \coloneqq \{\, e \mid \Gamma \vdash^{n+1} e : A \,\}$$

Function types are interpreted as dependent functions that, given a later-stage substitution from Γ to Γ' (where Γ' can be any later-stage context), and map elements of $(\!|\, A \,|\!)_{\Gamma', \Delta}$ to elements of $(\!|\, B \,|\!)_{\Gamma'}$. $\Delta \triangleright A$ is interpreted as the interpretation of A under the extended context Γ, Δ. **bool** is interpreted as the set of Booleans. $\bigcirc A$ is interpreted as the set of syntactic expressions of type A at level $n+1$, modulo α-equivalence.

We show that the interpretation is indeed monotonic. Given a type A, an element $d \in (\!|\, A \,|\!)_\Gamma$ (denoted d^A), and a later-stage substitution $\Gamma' \vdash \sigma : \Gamma$, we define $d[\sigma] \in (\!|\, A \,|\!)_{\Gamma'}$ as the result of applying the substitution σ to d. The definition proceeds by recursion on the type structure:

$\boxed{d[\sigma]}$ *(Element Substitution)*

$$f^{A \to B}[\sigma] \coloneqq \lambda \sigma'.\, f\ (\sigma[\sigma']) \qquad b^{\mathbf{bool}}[\sigma] \coloneqq b$$

$$d^{\Delta \triangleright A}[\sigma] \coloneqq d^A[\sigma, \mathsf{id}_\Delta] \qquad e^{\bigcirc A}[\sigma] \coloneqq e[\sigma]$$

7.2 Context Interpretation

Typing contexts at level n are interpreted as the product of the interpretations of their entries, where the interpretation of each entry depends on whether it is at the current stage n. Note that in

[6] This is similar to a *presheaf model* [Altenkirch et al. 2005; Kavvos 2024] over the category of later-stage substitutions, but without naturality conditions. Since our goal is to establish adequacy, which does not require naturality, we work with this simpler model and leave development of a presheaf model for future work.

the following definition, the context Γ being interpreted is at level n, while the context Γ' and Γ'' used as indices are at level $n+1$.

$\boxed{(\!|\,\Gamma\,|\!)_{\Gamma'}}$ *(Context Interpretation (Environments))*

$$(\!|\,\Gamma^n\,|\!)_{\Gamma'} \coloneqq \prod_{\Gamma \ni x:[\Delta \vdash^m A]} \begin{cases} \prod_{(\Gamma'' \vdash \sigma : \Gamma')} (\!|\,\Delta\,|\!)_{\Gamma''} \to (\!|\,A\,|\!)_{\Gamma''} & \text{if } m = n, \\ \{\, y \mid \Gamma' \ni y : [\Delta \vdash^m A] \,\} \cup \{\, e \mid \Gamma'\restriction_m, \Delta \vdash^m e : A \,\} & \text{if } m > n. \end{cases}$$

Specifically, current-stage entries $\Gamma \ni x : [\Delta \vdash^n A]$ are interpreted as dependent functions from $(\!|\,\Delta\,|\!)$ to $(\!|\,A\,|\!)$, similar to the interpretation of function types. Later-stage entries $\Gamma \ni x : [\Delta \vdash^m A]$ (where $m > n$) are interpreted as syntactic substitution entries under Γ', which can be either a variable $\Gamma' \ni y : [\Delta \vdash^m A]$ (rule S-Rename) or an expression $\Gamma'\restriction_m, \Delta \vdash^m e : A$ (rule S-Subst).

We call an element $\rho \in (\!|\,\Gamma\,|\!)_{\Gamma'}$ an *environment*, and write $\rho(x)$ to denote the entry corresponding to x in ρ. As with types, we show that the interpretation is monotonic. Given an environment $\rho \in (\!|\,\Gamma\,|\!)_{\Gamma'}$ and a later-stage substitution $\Gamma'' \vdash \sigma : \Gamma'$, we define $\rho[\sigma] \in (\!|\,\Gamma\,|\!)_{\Gamma''}$ by applying the substitution σ to each entry in ρ:

$\boxed{\rho[\sigma]}$ *(Environment Substitution)*

$$\rho[\sigma](x) \coloneqq \begin{cases} \lambda\sigma'.\ \rho(x)\ (\sigma[\sigma']) & \text{if } m = n, \\ e[\sigma\restriction_m, \mathrm{id}_\Delta] & \text{if } m > n \text{ and } \rho(x) = e, \\ \sigma(y) & \text{if } m > n \text{ and } \rho(x) = y, \end{cases} \quad \text{for each } \Gamma \ni x : [\Delta \vdash^m A].$$

where the case for $m = n$ is defined similarly to the function type, and the cases for $m > n$ are handled by substituting the substitution entry as defined in §4.1.

We write $\rho\restriction_{n+1}$ to remove all entries from ρ that are at level n, turning it into a substitution at level $n+1$. We can lift an element $d \in (\!|\,A\,|\!)_{\Gamma,\Delta}$ to a singleton environment $\{x^n \mapsto d\}$, whose definition is given as follows:

$$\{x^n \mapsto d\} \in \textstyle\prod_{\Gamma' \vdash \sigma : \Gamma} (\!|\,\Delta\,|\!)_{\Gamma'} \to (\!|\,A\,|\!)_{\Gamma'} \qquad \{x^n \mapsto d\} \coloneqq \lambda\sigma\ \rho.\ d[\sigma', \rho\restriction_{n+1}]$$

Note that in the above definition, since Δ is already at level $n+1$, $\rho\restriction_{n+1}$ does not remove any entry from ρ, but instead just converts it into a substitution. We write $\rho \cup \rho'$ to add entries to an environment, where ρ' can either be an environment or a later-stage substitution.

7.3 Expression Interpretation

Given a later-stage context Γ', expressions $\Gamma \vdash^n e : A$ are interpreted as functions $(\!|\,\Gamma\,|\!)_{\Gamma'} \to (\!|\,A\,|\!)_{\Gamma'}$, and substitutions $\Gamma \vdash \sigma : \Delta$ are interpreted as functions $(\!|\,\Gamma\,|\!)_{\Gamma'} \to (\!|\,\Delta\,|\!)_{\Gamma'}$. This interpretation can be viewed as a definitional interpreter for the language, which recursively interprets an expression under a given environment $\rho \in (\!|\,\Gamma\,|\!)_{\Gamma'}$ and extending the environment as needed. We present a selection of rules for expression interpretation for $\lambda^{\bigcirc\triangleright}$:

$\boxed{(\!|\,e\,|\!)_{\Gamma'}, (\!|\,\sigma\,|\!)_{\Gamma'}}$ *(Expression and Substitution Interpretation (excerpt))*

$$(\!|\,x_{\sigma_1}\,|\!)_{\Gamma'} \coloneqq \lambda\rho.\ \rho(x)\ \mathrm{id}_{\Gamma'}\ (\!|\,\sigma_1\,|\!)_{\Gamma'} \qquad (\!|\,\lambda_\Delta x : A.\ e\,|\!)_{\Gamma'} \coloneqq \lambda\rho.\ \lambda\sigma'\ d.\ (\!|\,e\,|\!)_{\Gamma'}\ (\rho[\sigma'] \cup \{x^n \mapsto d\})$$

$$(\!|\,\langle e \rangle\,|\!)_{\Gamma'} \coloneqq \lambda\rho.\ e[\rho\restriction_{n+1}] \qquad (\!|\,e_1\ e_2\,|\!)_{\Gamma'} \coloneqq \lambda\rho.\ (\!|\,e_1\,|\!)_{\Gamma'}\ \rho\ \mathrm{id}_{\Gamma'}\ ((\!|\,e_2\,|\!)_{\Gamma',\Delta}\ (\rho \cup \mathrm{id}_\Delta))$$

$$(\!|\,\mathbf{let}_\Delta \langle x : A \rangle = e_1\ \mathbf{in}\ e_2\,|\!)_{\Gamma'} \coloneqq \lambda\rho.\ \mathrm{let}\ e = (\!|\,e_1\,|\!)_{\Gamma',\Delta}\ (\rho \cup \mathrm{id}_\Delta)\ \mathrm{in}\ (\!|\,e_2\,|\!)_{\Gamma'}\ (\rho \cup (x \mapsto e))$$

For $\lambda^{\bigcirc\triangleright}_{\mathrm{pat}}$, the additional expression forms are interpreted as follows:

$$(\!|\,\mathbf{if}\ \mathbf{let}_\Delta \langle p \rangle = e_1\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3\,|\!)_{\Gamma'} \coloneqq \lambda\rho.$$

$$\begin{cases} (\!|\,e_2\,|\!)_{\Gamma'}\ (\rho \cup \sigma) & \text{if } \mathsf{match}\big(p[\rho\restriction_{n+1}];\ (\!|\,e_1\,|\!)_{\Gamma',\Delta}\ (\rho \cup \mathrm{id}_\Delta)\big) = \sigma, \\ (\!|\,e_3\,|\!)_{\Gamma'}\ \rho & \text{otherwise.} \end{cases}$$

$$(\!|\,e_1\ \mathbf{rewrite}\ \langle p_2 \rangle \mapsto e_2\,|\!)_{\Gamma'} \coloneqq \lambda\rho.\ \mathrm{rewrite}_{\mathsf{BottomUp}}(p_2[\rho\restriction_{n+1}];\ (\!|\,e_2\,|\!)_{\Gamma',\Pi}\ (\rho \cup \mathrm{id}_\Pi);\ (\!|\,e_1\,|\!)_{\Gamma'}\ \rho)$$

The rewrite interpretation corresponds to the small-step semantics (§6.3), where the replacement expression e_2 is evaluated before rewriting happens. Using the denotational semantics, we can alternatively evaluate the replacement expression at each occurrence of the pattern p within e_1, thus allowing the replacement expression to inspect the specific match results. This approach is analogous to the big-step semantics of the rewrite construct in Parreaux et al. [2017].

$$(\!| e_1 \textbf{ rewrite } \langle p_2 \rangle \mapsto e_2 |\!)_{\Gamma'} := \lambda\rho.\ \mathsf{applyBottomUp}\Big(\lambda e.\ \begin{cases} (\!| e_2 |\!)_{\Gamma'}\ (\rho \cup \sigma) & \text{if } p_2 \text{ and } e \text{ have the same type and } \mathsf{match}(p_2[\rho\restriction_{n+1}]; e) = \sigma, \\ e & \text{otherwise.} \end{cases} ; (\!| e_1 |\!)_{\Gamma'}\ \rho\Big)$$

This interpretation uses the applyBottomUp function to apply e_2's interpretation to every matching sub-expression of the target expression e_1. Crucially, the environment ρ is extended with the substitution σ that represents the match result, rather than id_Π as in the previous definition. This allows e_1 to inspect the match result using $\lambda^{\bigcirc\triangleright}_{\text{pat}}$'s analytic capabilities. The alternative semantics for rewriting is strictly more expressive than the small-step version, since the original semantics can be encoded in the alternative semantics by evaluating the replacement expression before rewriting.

7.4 Adequacy

To prove the adequacy of the denotational semantics, we define the logical relations $\mathcal{V}$, $\mathcal{E}$, and $\mathcal{G}$, which relate values, expressions, and environments with their denotation, respectively.

$\boxed{\mathcal{V}^A_\Gamma, \mathcal{E}^A_\Gamma}$ *(Value and Expression Relations)*

$$\mathcal{V}^{[\Delta \vdash A] \to B}_\Gamma := \{\ (\lambda_\Delta x : A.\ e; f) \mid \forall(\Gamma' \vdash \sigma : \Gamma).\ \forall(v; d) \in \mathcal{V}^A_{\Gamma',\Delta}.\ (e[\sigma, x \mapsto v]; f\ \sigma\ d) \in \mathcal{E}^B_{\Gamma'}\ \}$$

$$\mathcal{V}^{\Delta \triangleright A}_\Gamma := \{\ (\mathbf{wrap}_\Delta e; d) \mid (e; d) \in \mathcal{V}^A_{\Gamma,\Delta}\ \} \qquad \mathcal{V}^{\mathbf{bool}}_\Gamma := \{(\mathbf{true}; \mathsf{True}), (\mathbf{false}; \mathsf{False})\}$$

$$\mathcal{V}^{\bigcirc A}_\Gamma := \{\ (\langle e \rangle; e) \mid \Gamma \vdash^{n+1} e : A\ \}$$

$$\mathcal{E}^A_\Gamma := \{\ (e; d) \mid \exists v.\ (\Gamma \vdash^n e \longrightarrow^* v) \wedge (v; d) \in \mathcal{V}^A_\Gamma\ \}$$

$\boxed{\mathcal{G}^\Gamma_{\Gamma'}}$ *(Environment Relation)*

$$\mathcal{G}^\Gamma_{\Gamma'} := \prod_{\Gamma \ni x : [\Delta \vdash^m A]} \begin{cases} \{\ (e; f) \mid \forall(\Gamma'' \vdash \sigma : \Gamma').\ \forall(\sigma'; \rho') \in \mathcal{G}^\Delta_{\Gamma''}.\ (e[\sigma, \sigma']; f\ \sigma\ \rho') \in \mathcal{E}^A_{\Gamma''}\ \} & \text{if } m = n, \\ \{\ (y; y) \mid \Gamma' \ni y : [\Delta \vdash^m A]\ \} \cup \{\ (e; e) \mid \Gamma'\restriction_m, \Delta \vdash^m e : A\ \} & \text{if } m > n. \end{cases}$$

The relations generally follow the structure of the denotational semantics. The $\mathcal{V}^A_\Gamma$ relation relates values of type A under context Γ to elements in $(\!| A |\!)_\Gamma$. For base types **bool** and $\bigcirc A$, $\mathcal{V}$ is essentially the identity relation, relating each value to a unique element. For function types, a lambda value $(\lambda_\Delta x : A.\ e)$ is related to a function f if, for any later-stage substitution σ and any value v related to d, substituting e with σ and v yields an expression related to $f\ \sigma\ d$. For $\Delta \triangleright A$, the relation extends the context with Δ, as in the denotational semantics. The expression relation $\mathcal{E}$ relates an expression to an element if the expression reduces to a value related to that element. The environment relation $\mathcal{G}^\Gamma_{\Gamma'}$ relates substitutions with type $\Gamma' \vdash \sigma : \Gamma$ to elements in $(\!| \Gamma |\!)_{\Gamma'}$. For current-stage entries, an expression entry e is related to a function f in the same way lambda values are related to functions, while variable entries are not related. For later-stage entries, it relates the entry to itself.

We prove the fundamental property of the logical relations:

Theorem 7.1 (Fundamental Property).

- *If* $\Gamma \vdash^n e : A$ *and* $(\sigma; \rho) \in \mathcal{G}^\Gamma_{\Gamma'}$ *then* $(e[\sigma]; (\!| e |\!)_{\Gamma'}\ \rho) \in \mathcal{E}^A_{\Gamma'}$.
- *If* $\Gamma \vdash \sigma_1 : \Delta$ *and* $(\sigma; \rho) \in \mathcal{G}^\Gamma_{\Gamma'}$ *then* $(\sigma_1[\sigma]; (\!| \sigma_1 |\!)_{\Gamma'}\ \rho) \in \mathcal{G}^\Delta_{\Gamma'}$.

Table 1. Formalization in Agda

Definitions		$\lambda^{\bigcirc\triangleright}$	$\lambda^{\bigcirc\triangleright}_{\text{pat}}$
Data.StagedList/StagedTree	Intrinsically well-staged lists and rose trees.	314	
Context	Defines types and typing contexts.	22	22
Term	Defines intrinsically typed terms using de Bruijn indices.	142	262
Depth	Defines the depth of contexts and substitutions.	114	114
Substitution	Defines substitution.	34	77
Reduction	Defines operational semantics and proves safety properties.	95	87
Denotational	Defines the Kripke-style denotational semantics.	92	136
Examples	Examples of well-typed expressions and their evaluation.	66	67
Pattern Matching			
Matching	Defines the match function.		165
Rewriting	Defines the rewrite function.		77
Context.Equality	Proves decidable equality for $\lambda^{\bigcirc\triangleright}_{\text{pat}}$ contexts.		119
Term.Equality	Proves decidable equality for $\lambda^{\bigcirc\triangleright}_{\text{pat}}$ terms.		362
Adequacy			
Term.Properties	Proves properties of renaming operations.	175	
Substitution.Properties	Proves properties of substitution.	341	
Denotational.Adequacy	Proves the adequacy of the denotational semantics.	207	
Translation			
Context	Defines types and typing contexts in $\lambda^{\bigcirc}$.	30	
Term	Defines intrinsically typed terms in $\lambda^{\bigcirc}$.	68	
Translation	Formalizes the translation from $\lambda^{\bigcirc\triangleright}$ to $\lambda^{\bigcirc}$.	98	
Total		3286	

The fundamental property implies termination of the operational semantics and the adequacy of the denotational semantics. In particular, consider a level-n expression $\Gamma^{n+1} \vdash^n e : A$. Given $(\mathsf{id}_\Gamma; \mathsf{id}_\Gamma) \in \mathcal{G}^\Gamma_\Gamma$, it follows that $(e; (\!|\, e \,|\!)_\Gamma\ \mathsf{id}_\Gamma) \in \mathcal{E}^A_\Gamma$. Therefore, reducing e terminates at some value v. Furthermore, $\mathcal{V}$ relates v to $((\!|\, e \,|\!)_\Gamma\ \mathsf{id}_\Gamma)$, the result of evaluating e under the denotational semantics. Since the $\mathcal{V}$ is the identity relation at base types **bool** and $\bigcirc A$, we have $v = (\!|\, e \,|\!)_\Gamma\ \mathsf{id}_\Gamma$. That is, the result of the operational semantics coincides with the denotational semantics.

Theorem 7.2 (Adequacy).

- *If* $\Gamma^{n+1} \vdash^n e : \bigcirc A$, *then* $e \longrightarrow^* \langle e' \rangle$ *where* $e' = (\!|\, e \,|\!)_\Gamma\ \mathsf{id}_\Gamma$.
- *If* $\Gamma^{n+1} \vdash^n e :$ **bool**, *then* $e \longrightarrow^*$ **true** *if* $(\!|\, e \,|\!)_\Gamma\ \mathsf{id}_\Gamma = \mathsf{True}$, *and* $e \longrightarrow^*$ **false** *if* $(\!|\, e \,|\!)_\Gamma\ \mathsf{id}_\Gamma = \mathsf{False}$.

We have proven the fundamental property in Agda for $\lambda^{\bigcirc\triangleright}$, using an inductive proof on the typing derivation and employing the associativity and commutativity properties of substitution (§4.1.2). We have not yet fully established the fundamental property for $\lambda^{\bigcirc\triangleright}_{\text{pat}}$, but we expect it to hold as well. The main technical challenge is proving the substitution properties for $\lambda^{\bigcirc\triangleright}_{\text{pat}}$, which is more involved because the number of expression forms doubles with the introduction of patterns.

8 Formalization in Agda

We formalize the syntax, typing rules, semantics, translations, stated lemmas and theorems in Agda. Our formalization leverages the agda-stdlib library [The Agda Community 2024] and adopts the style of *Programming Language Foundations in Agda* [Wadler et al. 2022]. We check all proofs with the `safe` flag to ensure soundness. Table 1 presents the structure of our definitions and proofs.

In the formalization, all contexts and types are intrinsically well-staged, and all expressions are intrinsically typed. The modules Data.StagedList and Data.StagedTree define intrinsically well-staged lists and rose trees, respectively. These modules are developed in a self-contained manner,

Table 2. Comparison of our calculus with related work

	$\lambda^{\bigcirc\triangleright}, \lambda^{\bigcirc\triangleright}_{\text{pat}}$	$\lambda^{\bigcirc}$ [Davies 1996]	$F^{[\![]\!]}$ [Xie et al. 2022]	Mœbius [Jang et al. 2022]	$\lambda^{\blacktriangle}$ [Stucki et al. 2021]	$\lambda^{\{\}}$ [Parreaux et al. 2017]
Quoting Unquoting	$\langle\cdot\rangle$ $\mathbf{let}_{\Delta}\ \langle\cdot\rangle$	**next** **prev**	$[\![\cdot]\!]_{\phi}$	box let box	$\lceil\cdot\rceil$ $\lfloor\cdot\rfloor$	$\lceil\cdot\rceil$ $\lfloor\cdot\rfloor$
Code Type Contextual Type	$\bigcirc$ $\Delta \triangleright$	$\bigcirc$ –	*Code* –	$\lceil \Phi \vdash^{k} \cdot \rceil$	$\lceil\cdot\rceil$ –	**Code** $T\ C$
Nested Contexts	Yes	–	1-level	Yes	–	No
Polymorphism	No	No	Yes	Yes	No	Subtyping
Pattern Matching Rewrite	Yes Yes	– –	– –	Yes –	Yes –	Yes Yes

making them reusable for other projects requiring well-staged data structures. In our formalization, they are used to represent the nested structure of typing contexts. In Data.StagedList, concatenation of well-staged lists is defined as a constructor rather than a function, effectively making it a tree rather than a list. This design choice allows other context operations to commute with concatenation by definition, simplifying proofs by reducing the need for casting between equivalent contexts. Variables are represented namelessly using de Bruijn indices. These simplifications contribute to a more concise formalization and ensure that pattern matching respects α-equivalence.

9 Related Work

This section compares our calculi with related work. Table 2 provides an overview of the syntax, type systems, and key features of these calculi.

Let-splice bindings. Let-splice bindings are inspired by $F^{[\![]\!]}$, the core calculus of Typed Template Haskell [Xie et al. 2022], but with several key distinctions. In $F^{[\![]\!]}$, let-splice bindings (called *splice environments*) are bound to quotations, appearing as $[\![e]\!]_{\phi}$, where e is a quoted expression and ϕ is a list of splice bindings. This effectively ties splice bindings to quotations (except for top-level splices). Our calculus decouples let-splice bindings from quotations, allowing for greater flexibility. Moreover, $F^{[\![]\!]}$ contexts are flat, supporting only a single level of nesting, and does not support dependencies ($\Delta \triangleright$) as first-class types. Our calculus allows arbitrary nesting of contexts, enabling more complex variable dependencies. This nesting necessitates delayed substitutions to ensure type soundness, as discussed in §4.2. Finally, $F^{[\![]\!]}$ lacks support for code pattern matching. In contrast to syntax-directed staging, Scala's LMS (Lightweight Modular Staging) framework [Rompf and Odersky 2010, 2012] mixes present and future-stage computations using a type-based embedding, reducing syntactic overhead. It would be interesting to explore how concepts similar to let-splice bindings could potentially apply to such a type-directed staging approach.

Modal logic and staging. While temporal logic based staging [Davies 1996] has been used widely, S4 modal logic offers an alternative approach to staging [Pfenning and Davies 2001]. Specifically, the box modality $\Box A$ models *closed code expressions*, which cannot depend on the surrounding context. Thus, $\Box A$ can always be evaluated to yield a value of type A. This contrasts with the temporal modality $\bigcirc A$, which allows code to use variables in the surrounding context in a well-staged way. The relationship between $\lambda^{\bigcirc}$, $F^{[\![]\!]}$, and $\lambda^{\bigcirc\triangleright}$ mirrors the different derivation systems of the intuitionistic S4 logic. $\lambda^{\bigcirc}$ corresponds to Pfenning and Davies [2001]'s implicit system which uses quote and unquote operators similar to quasi-quotes, $F^{[\![]\!]}$ corresponds to the style of Bierman and de Paiva [2000], pairing an explicit substitution with the quote constructor, and $\lambda^{\bigcirc\triangleright}$ corresponds to Pfenning and Davies [2001]'s explicit system, using let-bindings for unquoting. In literature, implicit

systems are sometimes referred to as Kripke-style or Fitch-style [Clouston 2018; Murase 2017; Murase et al. 2023], while explicit systems are sometimes called dual-context style [Kavvos 2020; Nanevski et al. 2008]. *Contextual modal type theory* (CMTT) [Nanevski et al. 2008] extends the S4 approach with *contextual modalities*, generalizing the $\Box A$ type to $\lceil \Phi \vdash A \rceil$, which explicitly declares a context Ψ that the code may reference. Code of type $\lceil \cdot \vdash A \rceil$ is guaranteed to be closed (containing no free variables), and can thus be evaluated. Our substitution for $x : [\Delta \vdash^n A]$ is similar to meta-substitution in second-order abstract syntax [Fiore and Hur 2010], as well as to CMTT. Boespflug and Pientka [2011] further extend CMTT with multiple levels, modeling meta^n-variables. Mœbius [Jang et al. 2022] additionally supports code pattern matching. Our calculus shares similarities with Mœbius, but with important differences. First, we have different logical foundations: ours is based on temporal logic, while Mœbius generalizes from S4. Thus, our $\bigcirc A$ can still access variables from its surrounding context. Within a quotation, our system treats both meta-variables and program variables as variables at the next level, while in Mœbius, meta-variables are treated separately from program variables. Also, we separate the code modality $\bigcirc$ and the contextual modality $(\Delta \rhd)$, while Mœbius combines them into a single modality $\lceil \Phi \vdash^k \cdot \rceil$. Furthermore, we use different approaches to context tracking: a value of type $(\Delta \rhd A)$ can use variables in Δ *in addition to* those in the surrounding context. This allows dependencies that do not follow lexical scoping, which is essential for expressing unhygienic functions. It also allows context specifications to be combined with other type constructors. For instance, $(x : \mathbf{bool}^1) \rhd (A \times \bigcirc B)$ could represent an unhygienic value of type A paired with code of type B, where both use a variable x. This different logical foundation and design also allow us to use different types for pattern variables. Lastly, Mœbius supports type polymorphism; as future work, we are interested in extending our design with type polymorphism.

Multimodal type theory [Gratzer et al. 2020; Kavvos and Gratzer 2023] provides a general framework for combining multiple modal types in a single type system. It seems possible to consider $\lambda^{\bigcirc\rhd}$ as a multimodal type theory with integer modes with $\bigcirc : n + 1 \to n$, whose 🔒 corresponds to context restriction, and $\Delta^{n+1} \rhd : n \to n$, whose 🔒 corresponds to context extension. The rules for unhygienic functions can be seen as a special case of function types in multimodal type theory, where the modality is fixed to be $\Delta^{n+1} \rhd$. Specifying our calculus as a multimodal system would be an interesting direction.

Code pattern matching. Stucki et al. [2021] present $\lambda^{\blacktriangle}$, which also supports code pattern matching. In $\lambda^{\blacktriangle}$, as mentioned in §2.5, to match a term under a lambda, the pattern variable must first be η-expanded into a function that takes the code of the bindings as arguments. As noted by Stucki et al., this only works for a simpler two-stage setting and does not yet support matching on quotes, splices, and matches themselves. Compared to $\lambda^{\blacktriangle}$'s approach, our type system avoids the need for η-expansion and allows for matching on, for example, code that itself uses quotations and splices. Squid [Parreaux et al. 2017] (and their calculus $\lambda^{\{\}}$) uses CMTT-like types and is thus closely related to Mœbius [Jang et al. 2022]. Our code rewriting construct is inspired by Squid.

Gabbay and Nanevski [2013] attempt to provide a denotational semantics for CMTT, which includes an intensional isapp function that checks whether a code is a function application. As noted by Kavvos [2017a,b], their evaluation of isapp is not confluent. Kavvos [2017a,b] extend the modal lambda calculus with intensional analysis on closed code to retain confluence. Hu and Pientka [2024, 2025] propose *layered modal type theory* (LMTT), where the language at layer $i + 1$ can inspect and analyze code from the one at layer i. Similar to CMTT, LMTT distinguishes between local variables and global variables, and additionally also between code and programs (i.e. layers), while our calculus uses the same language for code and programs, and also the same variable dependency mechanism for splice variables and pattern variables. We intend to explore confluence in future work.

Acknowledgments

This work is funded by the Natural Sciences and Engineering Research Council of Canada.

Data Availability Statement

The Agda mechanization of this paper is provided in the artifact [Chiang and Xie 2025].

A $\lambda^{\bigcirc\triangleright}$ Details

A.1 Syntax

$$
\begin{array}{llll}
\text{Variables} & x, y, z & & \\
\text{Levels} & m, n \in \mathbb{N} & & \\
\text{Types} & A, B & ::= & \mathbf{bool} \mid [\Delta \vdash A] \to B \mid \bigcirc A \mid \Delta \triangleright A \\
\text{Contexts} & \Gamma, \Delta & ::= & \cdot \mid \Gamma, x : [\Delta \vdash^n A] \\
\text{Expressions} & e & ::= & x_\sigma \mid \mathbf{true} \mid \mathbf{false} \mid \mathbf{if}\ e_1\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3 \\
 & & \mid & \lambda_\Delta x : A.\ e \mid e_1\ e_2 \mid \langle e \rangle \mid \mathbf{let}_\Delta \langle x : A \rangle = e_1\ \mathbf{in}\ e_2 \\
 & & \mid & \mathbf{wrap}_\Delta e \mid \mathbf{let}\,\mathbf{wrap}_\Delta\ x : A = e_1\ \mathbf{in}\ e_2 \\
\text{Substitutions} & \sigma & ::= & \cdot \mid \sigma, x \mapsto y \mid \sigma, x \mapsto e
\end{array}
$$

Fig. 3. Syntax of $\lambda^{\bigcirc\triangleright}$

A.2 Substitution

$\boxed{e[\sigma]}$ *(Expression Substitution)*

$$
\begin{aligned}
(x_{\sigma_1})[\sigma] &\coloneqq \begin{cases} y_{(\sigma_1[\sigma])} & \text{if } \sigma(x) = y, \\ e[\mathsf{id}_{\Gamma_2 \restriction_m}, \sigma_1[\sigma]] & \text{if } \sigma(x) = e. \end{cases} \\
(\mathbf{true})[\sigma] &\coloneqq \mathbf{true} \\
(\mathbf{false})[\sigma] &\coloneqq \mathbf{false} \\
(\mathbf{if}\ e_1\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3)[\sigma] &\coloneqq \mathbf{if}\ e_1[\sigma]\ \mathbf{then}\ e_2[\sigma]\ \mathbf{else}\ e_3[\sigma] \\
(\lambda_\Delta x : A.\ e)[\sigma] &\coloneqq \lambda_\Delta x : A.\ (e[\sigma, x \mapsto x]) \\
(e_1\ e_2)[\sigma] &\coloneqq e_1[\sigma]\ e_2[\sigma, \mathsf{id}_\Delta] \\
(\langle e \rangle)[\sigma] &\coloneqq \langle e[\sigma\restriction] \rangle \\
(\mathbf{let}_\Delta \langle x : A \rangle = e_1\ \mathbf{in}\ e_2)[\sigma] &\coloneqq \mathbf{let}_\Delta \langle x : A \rangle = (e_1[\sigma, \mathsf{id}_\Delta])\ \mathbf{in}\ (e_2[\sigma, x \mapsto x]) \\
(\mathbf{wrap}_\Delta e)[\sigma] &\coloneqq \mathbf{wrap}_\Delta (e[\sigma, \mathsf{id}_\Delta]) \\
(\mathbf{let}\,\mathbf{wrap}_\Delta\ x : A = e_1\ \mathbf{in}\ e_2)[\sigma] &\coloneqq \mathbf{let}\,\mathbf{wrap}_\Delta\ x : A = (e_1[\sigma])\ \mathbf{in}\ (e_2[\sigma, x \mapsto x])
\end{aligned}
$$

$\boxed{\sigma_1[\sigma]}$ *(Substitution on Substitutions)*

$$
\begin{aligned}
(\cdot)[\sigma] &\coloneqq \cdot & (\sigma_1, x \mapsto y)[\sigma] &\coloneqq \sigma_1[\sigma], x \mapsto \sigma(y) \\
 & & (\sigma_1, x^m_\Delta \mapsto e)[\sigma] &\coloneqq \sigma_1[\sigma], x \mapsto e[\sigma\restriction_m, \mathsf{id}_\Delta]
\end{aligned}
$$

A.3 Operational Semantics

$$
\begin{array}{lll}
\text{Values} & v & ::= \mathbf{true} \mid \mathbf{false} \mid \lambda_\Delta x : A.\ e \mid \langle e \rangle \mid \mathbf{wrap}_\Delta v \\
\text{Evaluation Contexts} & E & ::= [\,] \mid E\ e_2 \mid v_1\ E \mid \mathbf{if}\ E\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3 \mid \mathbf{let}_\Delta \langle x : A \rangle = E\ \mathbf{in}\ e_2 \\
 & & \mid \mathbf{wrap}_\Delta E \mid \mathbf{let}\,\mathbf{wrap}_\Delta\ x : A = E\ \mathbf{in}\ e_2
\end{array}
$$

$\boxed{\Gamma \vdash e \longrightarrow e'}$ *(Call-by-value Reduction)*

$$\textsc{CtxAppAbs}\ \frac{}{\Gamma \vdash (\lambda_\Delta x : A.\ e_1)\ v_2 \longrightarrow e_1[\mathrm{id}_\Gamma, x \mapsto v_2]} \qquad \textsc{LetQuoteQuote}\ \frac{}{\Gamma \vdash \mathbf{let}_\Delta \langle x : A \rangle = \langle e_1 \rangle \mathbf{\ in\ } e_2 \longrightarrow e_2[\mathrm{id}_\Gamma, x \mapsto e_1]}$$

$$\textsc{IfTrue}\ \frac{}{\Gamma \vdash \mathbf{if\ true\ then\ } e_2 \mathbf{\ else\ } e_3 \longrightarrow e_2} \qquad \textsc{IfFalse}\ \frac{}{\Gamma \vdash \mathbf{if\ false\ then\ } e_2 \mathbf{\ else\ } e_3 \longrightarrow e_3}$$

$$\textsc{LetWrapWrap}\ \frac{}{\Gamma \vdash \mathbf{let\ wrap}_\Delta\ x : A = \mathbf{wrap}_\Delta\ v_1 \mathbf{\ in\ } e_2 \longrightarrow e_2[\mathrm{id}_\Gamma, x \mapsto v_1]} \qquad \textsc{CongCtxApp1}\ \frac{\Gamma \vdash e_1 \longrightarrow e_1'}{\Gamma \vdash e_1\ e_2 \longrightarrow e_1'\ e_2}$$

$$\textsc{CongCtxApp2}\ \frac{\Gamma, \Delta \vdash e_2 \longrightarrow e_2'}{\Gamma \vdash (\lambda_\Delta x : A.\ e_1)\ e_2 \longrightarrow (\lambda_\Delta x : A.\ e_1)\ e_2'} \qquad \textsc{CongLetQuote1}\ \frac{\Gamma, \Delta \vdash e_1 \longrightarrow e_1'}{\Gamma \vdash \mathbf{let}_\Delta \langle x : A \rangle = e_1 \mathbf{\ in\ } e_2 \longrightarrow \mathbf{let}_\Delta \langle x : A \rangle = e_1' \mathbf{\ in\ } e_2}$$

$$\textsc{CongIf1}\ \frac{\Gamma \vdash e_1 \longrightarrow e_1'}{\Gamma \vdash \mathbf{if}\ e_1 \mathbf{\ then\ } e_2 \mathbf{\ else\ } e_3 \longrightarrow \mathbf{if}\ e_1' \mathbf{\ then\ } e_2 \mathbf{\ else\ } e_3} \qquad \textsc{CongWrap}\ \frac{\Gamma, \Delta \vdash e \longrightarrow e'}{\Gamma \vdash \mathbf{wrap}_\Delta e \longrightarrow \mathbf{wrap}_\Delta e'}$$

$$\textsc{CongLetWrap1}\ \frac{\Gamma \vdash e_1 \longrightarrow e_1'}{\Gamma \vdash \mathbf{let\ wrap}_\Delta\ x : A = e_1 \mathbf{\ in\ } e_2 \longrightarrow \mathbf{let\ wrap}_\Delta\ x : A = e_1' \mathbf{\ in\ } e_2}$$

B $\lambda^{\bigcirc\triangleright}_{\text{pat}}$ Details

B.1 Syntax

Expressions	e	::=	$\ldots \mid \mathbf{if\ let}_\Delta \langle p \rangle = e_1 \mathbf{\ then\ } e_2 \mathbf{\ else\ } e_3 \mid e_1 \mathbf{\ rewrite\ } \langle p \rangle \mapsto e_2$
Patterns	p	::=	$\hat{x} : A \mid x_\pi \mid \mathbf{true} \mid \mathbf{false} \mid \mathbf{if}\ p_1 \mathbf{\ then\ } p_2 \mathbf{\ else\ } p_3$
		$\mid$	$\lambda_\Delta x : A.\ p \mid p_1\ p_2 \mid \langle p \rangle \mid \mathbf{let}_\Delta \langle x : A \rangle = p_1 \mathbf{\ in\ } p_2$
		$\mid$	$\mathbf{wrap}_\Delta p \mid \mathbf{let\ wrap}_\Delta\ x : A = p_1 \mathbf{\ in\ } p_2$
		$\mid$	$\mathbf{if\ let}_\Delta \langle p \rangle = p_1 \mathbf{\ then\ } p_2 \mathbf{\ else\ } p_3 \mid p_1 \mathbf{\ rewrite\ } \langle p \rangle \mapsto p_2$
Substitution Patterns	π	::=	$\cdot \mid \pi, x \mapsto y \mid \pi, x \mapsto p$

Fig. 4. Additional syntax forms of $\lambda^{\bigcirc\triangleright}_{\text{pat}}$

B.2 Additional Typing Rules

$\boxed{\Gamma \vdash^n e : A}$ *(Expression Typing)*

$$\textsc{IfLet}\ \frac{\Gamma\restriction_{n+1}; \Delta^{n+1} \vdash^{n+1} p : A \rightsquigarrow \Pi^{n+1} \qquad \Gamma, \Delta \vdash^n e_1 : \bigcirc A \qquad \Gamma, \Pi \vdash^n e_2 : B \qquad \Gamma \vdash^n e_3 : B}{\Gamma \vdash^n \mathbf{if\ let}_\Delta \langle p \rangle = e_1 \mathbf{\ then\ } e_2 \mathbf{\ else\ } e_3 : B}$$

$$\textsc{Rewrite}\ \frac{\Gamma \vdash^n e_1 : \bigcirc A \qquad \Gamma\restriction_{n+1}; \cdot \vdash^{n+1} p : B \rightsquigarrow \Pi^{n+1} \qquad \Gamma, \Pi \vdash^n e_2 : \bigcirc B}{\Gamma \vdash^n e_1 \mathbf{\ rewrite\ } \langle p \rangle \mapsto e_2 : \bigcirc A}$$

$\boxed{\Gamma; \Delta \vdash^n p : A \rightsquigarrow \Pi}$ *(Pattern Typing)*

$$\textsc{P-PVar}\ \frac{}{\Gamma; \Delta \vdash^n (\hat{x} : A) : A \rightsquigarrow x : [\Delta \vdash^n A]} \qquad \textsc{P-VarSubst1}\ \frac{\Gamma \ni x : [\Delta' \vdash^n A] \qquad \Gamma, \Delta \vdash \sigma : \Delta'}{\Gamma; \Delta \vdash^n x_\sigma : A \rightsquigarrow \cdot}$$

$$\textsc{P-VarSubst2}\ \frac{\Delta \ni x : [\Delta' \vdash^n A] \qquad \Gamma; \Delta \vdash \pi : \Delta' \rightsquigarrow \Pi}{\Gamma; \Delta \vdash^n x_\pi : A \rightsquigarrow \Pi}$$

$$\text{P-If}\ \frac{\Gamma;\Delta \vdash^n p_1 : \mathbf{bool} \rightsquigarrow \Pi_1 \qquad \Gamma;\Delta \vdash^n p_2 : A \rightsquigarrow \Pi_2 \qquad \Gamma;\Delta \vdash^n p_3 : A \rightsquigarrow \Pi_3}{\Gamma;\Delta \vdash^n \mathbf{if}\ p_1\ \mathbf{then}\ p_2\ \mathbf{else}\ p_3 : A \rightsquigarrow \Pi_1, \Pi_2, \Pi_3} \qquad \text{P-True}\ \frac{}{\Gamma;\Delta \vdash^n \mathbf{true} : \mathbf{bool} \rightsquigarrow \cdot}$$

$$\text{P-False}\ \frac{}{\Gamma;\Delta \vdash^n \mathbf{false} : \mathbf{bool} \rightsquigarrow \cdot} \qquad \text{P-CtxAbs}\ \frac{\Gamma;\Delta, x : [\Delta'^{n+1} \vdash^n A] \vdash^n p : B \rightsquigarrow \Pi}{\Gamma;\Delta \vdash^n (\lambda_{\Delta'} x : A.\ p) : [\Delta' \vdash A] \to B \rightsquigarrow \Pi}$$

$$\text{P-CtxApp}\ \frac{\Gamma;\Delta \vdash^n p_1 : [\Delta' \vdash A] \to B \rightsquigarrow \Pi_1 \qquad \Gamma;\Delta, \Delta' \vdash^n p_2 : A \rightsquigarrow \Pi_2}{\Gamma;\Delta \vdash^n p_1\ p_2 : B \rightsquigarrow \Pi_1, \Pi_2} \qquad \text{P-Quote}\ \frac{\Gamma\restriction_{n+1};\Delta\restriction_{n+1} \vdash^{n+1} p : A \rightsquigarrow \Pi^{n+1}}{\Gamma;\Delta \vdash^n \langle p \rangle : \bigcirc A \rightsquigarrow \Pi}$$

$$\text{P-LetQuote}\ \frac{\Gamma;\Delta, \Delta' \vdash^n p_1 : \bigcirc A \rightsquigarrow \Pi_1 \qquad \Gamma;\Delta, x : [\Delta' \vdash^{n+1} A] \vdash^n p_2 : B \rightsquigarrow \Pi_2}{\Gamma;\Delta \vdash^n \mathbf{let}_{\Delta'} \langle x : A \rangle = p_1\ \mathbf{in}\ p_2 : B \rightsquigarrow \Pi_1, \Pi_2} \qquad \text{P-Wrap}\ \frac{\Gamma;\Delta, \Delta' \vdash^n p : A \rightsquigarrow \Pi}{\Gamma;\Delta \vdash^n \mathbf{wrap}_{\Delta'} p : \Delta' \rhd A \rightsquigarrow \Pi}$$

$$\text{P-LetWrap}\ \frac{\Gamma;\Delta \vdash^n p_1 : \Delta' \rhd A \rightsquigarrow \Pi_1 \qquad \Gamma;\Delta, x : [\Delta' \vdash^n A] \vdash^n p_2 : B \rightsquigarrow \Pi_2}{\Gamma;\Delta \vdash^n \mathbf{let}\ \mathbf{wrap}_{\Delta'}\ x : A = p_1\ \mathbf{in}\ p_2 : B \rightsquigarrow \Pi_1, \Pi_2}$$

$$\text{P-IfLet}\ \frac{\Gamma;\Delta, \Delta' \vdash^n p_1 : \bigcirc A \rightsquigarrow \Pi_1 \qquad (\Gamma,\Delta)\restriction_{n+1};\Delta' \vdash^{n+1} p : A \rightsquigarrow \Pi \qquad \Gamma;\Delta, \Pi \vdash^n p_2 : B \rightsquigarrow \Pi_2 \qquad \Gamma;\Delta \vdash^n p_3 : B \rightsquigarrow \Pi_3}{\Gamma;\Delta \vdash^n \mathbf{if}\ \mathbf{let}_{\Delta'} \langle p \rangle = p_1\ \mathbf{then}\ p_2\ \mathbf{else}\ p_3 : B \rightsquigarrow \Pi_1, \Pi_2, \Pi_3}$$

$$\text{P-Rewrite}\ \frac{\Gamma;\Delta \vdash^n p_1 : \bigcirc A \rightsquigarrow \Pi_1 \qquad (\Gamma,\Delta)\restriction_{n+1};\cdot \vdash^{n+1} p : B \rightsquigarrow \Pi \qquad \Gamma;\Delta, \Pi \vdash^n p_2 : \bigcirc B \rightsquigarrow \Pi_2}{\Gamma;\Delta \vdash^n p_1\ \mathbf{rewrite}\ \langle p \rangle \mapsto p_2 : \bigcirc A \rightsquigarrow \Pi_1, \Pi_2}$$

$\boxed{\Gamma;\Delta \vdash \pi : \Gamma' \rightsquigarrow \Pi}$ *(Substitution Pattern Typing)*

$$\text{P-S-Empty}\ \frac{}{\Gamma;\Delta \vdash \cdot : \cdot \rightsquigarrow \cdot} \qquad \text{P-S-Var}\ \frac{\Gamma;\Delta \vdash \pi : \Gamma' \rightsquigarrow \Pi \qquad \Gamma,\Delta \ni y : [\Delta' \vdash^m A]}{\Gamma;\Delta \vdash (\pi, x \mapsto y) : \Gamma', x : [\Delta' \vdash^m A] \rightsquigarrow \Pi}$$

$$\text{P-S-Pattern}\ \frac{\Gamma;\Delta \vdash \pi : \Gamma' \rightsquigarrow \Pi_1 \qquad \Gamma\restriction_m;\Delta\restriction_m, \Delta'^m \vdash^m p : A \rightsquigarrow \Pi_2}{\Gamma;\Delta \vdash (\pi, x \mapsto p) : \Gamma', x : [\Delta' \vdash^m A] \rightsquigarrow \Pi_1, \Pi_2}$$

B.3 Operational Semantics

Values	v	$::= \mathbf{true} \mid \mathbf{false} \mid \lambda_\Delta x : A.\ e \mid \langle e \rangle \mid \mathbf{wrap}_\Delta v$
Evaluation Contexts	E	$::= [] \mid E\ e_2 \mid v_1\ E \mid \mathbf{if}\ E\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3 \mid \mathbf{let}_\Delta \langle x : A \rangle = E\ \mathbf{in}\ e_2$
		$\mid \mathbf{wrap}_\Delta E \mid \mathbf{let}\ \mathbf{wrap}_\Delta\ x : A = E\ \mathbf{in}\ e_2 \mid \mathbf{if}\ \mathbf{let}_\Delta \langle p \rangle = E\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3$
		$\mid E\ \mathbf{rewrite}\ \langle p \rangle \mapsto e_2 \mid v_1\ \mathbf{rewrite}\ \langle p \rangle \mapsto E$

$\boxed{\Gamma \vdash e \longrightarrow e'}$ *(Additional Reduction Rules)*

$$\text{IfLetQuote1}\ \frac{\mathsf{match}(p; e_1) = \sigma}{\Gamma \vdash \mathbf{if}\ \mathbf{let}_\Delta \langle p \rangle = \langle e_1 \rangle\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3 \longrightarrow e_2[\mathsf{id}_\Gamma, \sigma]} \qquad \text{IfLetQuote2}\ \frac{\mathsf{match}(p; e_1)\ \mathbf{undefined}}{\Gamma \vdash \mathbf{if}\ \mathbf{let}_\Delta \langle p \rangle = \langle e_1 \rangle\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3 \longrightarrow e_3}$$

$$\text{RewriteQuoteQuote}\ \frac{}{\Gamma \vdash \langle e_1 \rangle\ \mathbf{rewrite}\ \langle p \rangle \mapsto \langle e_2 \rangle \longrightarrow \langle \mathsf{rewrite}_{\mathsf{BottomUp}}(e_1; p; e_2) \rangle}$$

$$\textsc{CongIfLet1}\quad \frac{\Gamma, \Delta \vdash e_1 \longrightarrow e_1'}{\Gamma \vdash \textbf{if let}_\Delta \langle p \rangle = e_1 \textbf{ then } e_2 \textbf{ else } e_3 \longrightarrow \textbf{if let}_\Delta \langle p \rangle = e_1' \textbf{ then } e_2 \textbf{ else } e_3}$$

$$\textsc{CongRewrite1}\quad \frac{\Gamma \vdash e_1 \longrightarrow e_1'}{\Gamma \vdash e_1 \textbf{ rewrite } \langle p \rangle \mapsto e_2 \longrightarrow e_1' \textbf{ rewrite } \langle p \rangle \mapsto e_2}$$

$$\textsc{CongRewrite2}\quad \frac{\Gamma; \cdot \vdash^{n+1} p : B \rightsquigarrow \Pi \qquad \Gamma, \Pi \vdash e_2 \longrightarrow e_2'}{\Gamma \vdash \langle e_1 \rangle \textbf{ rewrite } \langle p \rangle \mapsto e_2 \longrightarrow \langle e_1 \rangle \textbf{ rewrite } \langle p \rangle \mapsto e_2'}$$

B.4 Pattern Matching

$\boxed{\mathsf{match}(p; e)}$ *(Expression Matching)*

$$\begin{aligned}
\mathsf{match}(\hat{x} : A; e) &\coloneqq x \mapsto e \\
\mathsf{match}(x_\sigma; x_\sigma) &\coloneqq \cdot \\
\mathsf{match}(x_\pi; x_\sigma) &\coloneqq \mathsf{match}(\pi; \sigma) \\
\mathsf{match}(\textbf{true}; \textbf{true}) &\coloneqq \cdot \\
\mathsf{match}(\textbf{false}; \textbf{false}) &\coloneqq \cdot \\
\mathsf{match}(\textbf{if } p_1 \textbf{ then } p_2 \textbf{ else } p_3; & \\
\qquad \textbf{if } e_1 \textbf{ then } e_2 \textbf{ else } e_3) &\coloneqq \mathsf{match}(p_1; e_1), \mathsf{match}(p_2; e_2), \mathsf{match}(p_3; e_3) \\
\mathsf{match}((\lambda_\Delta x : A.\ p); (\lambda_\Delta x : A.\ e)) &\coloneqq \mathsf{match}(p; e) \\
\mathsf{match}(p_1\ p_2; e_1\ e_2) &\coloneqq \mathsf{match}(p_1; e_1), \mathsf{match}(p_2; e_2) \\
\mathsf{match}(\langle p \rangle; \langle e \rangle) &\coloneqq \mathsf{match}(p; e) \\
\mathsf{match}(\textbf{let}_\Delta \langle x : A \rangle = p_1 \textbf{ in } p_2; & \\
\qquad \textbf{let}_\Delta \langle x : A \rangle = e_1 \textbf{ in } e_2) &\coloneqq \mathsf{match}(p_1; e_1), \mathsf{match}(p_2; e_2) \\
\mathsf{match}(\textbf{wrap}_\Delta p; \textbf{wrap}_\Delta e) &\coloneqq \mathsf{match}(p; e) \\
\mathsf{match}(\textbf{let wrap}_\Delta\ x : A = p_1 \textbf{ in } p_2; & \\
\qquad \textbf{let wrap}_\Delta\ x : A = e_1 \textbf{ in } e_2) &\coloneqq \mathsf{match}(p_1; e_1), \mathsf{match}(p_2; e_2) \\
\mathsf{match}(\textbf{if let}_\Delta \langle p \rangle = p_1 \textbf{ then } p_2 \textbf{ else } p_3; & \\
\qquad \textbf{if let}_\Delta \langle p \rangle = e_1 \textbf{ then } e_2 \textbf{ else } e_3) &\coloneqq \mathsf{match}(p_1; e_1), \mathsf{match}(p_2; e_2), \mathsf{match}(p_3; e_3) \\
\mathsf{match}(p_1 \textbf{ rewrite } \langle p \rangle \mapsto p_2; & \\
\qquad e_1 \textbf{ rewrite } \langle p \rangle \mapsto e_2) &\coloneqq \mathsf{match}(p_1; e_1), \mathsf{match}(p_2; e_2)
\end{aligned}$$

$\boxed{\mathsf{match}(\pi; \sigma)}$ *(Substitution Matching)*

$$\begin{aligned}
\mathsf{match}(\cdot; \cdot) &\coloneqq \cdot \\
\mathsf{match}(\pi, x \mapsto y; \sigma, x \mapsto y) &\coloneqq \mathsf{match}(\pi; \sigma) \\
\mathsf{match}(\pi, x \mapsto p; \sigma, x \mapsto e) &\coloneqq \mathsf{match}(\pi; \sigma), \mathsf{match}(p; e)
\end{aligned}$$

B.5 Rewriting Strategies

The meta-level function applyBottomUp(F, e) is defined as follows, where n is the level of e. For brevity, applyBottomUp(F, e) is denoted as $\mathsf{bu}(F; e)$.

$\boxed{\mathsf{bu}(F; e)}$ *(Bottom-up Application)*

$$\begin{aligned}
\mathsf{bu}(F; x_\sigma) &\coloneqq F(x_{\mathsf{bu}(F;\sigma)}) \\
\mathsf{bu}(F; \textbf{true}) &\coloneqq F(\textbf{true}) \\
\mathsf{bu}(F; \textbf{false}) &\coloneqq F(\textbf{false}) \\
\mathsf{bu}(F; \textbf{if } e_1 \textbf{ then } e_2 \textbf{ else } e_3) &\coloneqq F(\textbf{if } \mathsf{bu}(F; e_1) \textbf{ then } \mathsf{bu}(F; e_2) \textbf{ else } \mathsf{bu}(F; e_3)) \\
\mathsf{bu}(F; (\lambda_\Delta x : A.\ e)) &\coloneqq F(\lambda_\Delta x : A.\ \mathsf{bu}(F; e))
\end{aligned}$$

$$
\begin{aligned}
\mathsf{bu}(F;\, e_1\ e_2) &\coloneqq F(\mathsf{bu}(F;\, e_1)\ \mathsf{bu}(F;\, e_2)) \\
\mathsf{bu}(F;\, \langle e \rangle) &\coloneqq F(\langle e \rangle) \\
\mathsf{bu}(F;\, \mathbf{let}_\Delta \langle x : A \rangle = e_1\ \mathbf{in}\ e_2) &\coloneqq F(\mathbf{let}_\Delta \langle x : A \rangle = \mathsf{bu}(F;\, e_1)\ \mathbf{in}\ \mathsf{bu}(F;\, e_2)) \\
\mathsf{bu}(F;\, \mathbf{wrap}_\Delta e) &\coloneqq F(\mathbf{wrap}_\Delta \mathsf{bu}(F;\, e)) \\
\mathsf{bu}(F;\, \mathbf{let\ wrap}_\Delta\ x : A = e_1\ \mathbf{in}\ e_2) &\coloneqq F(\mathbf{let\ wrap}_\Delta\ x : A = \mathsf{bu}(F;\, e_1)\ \mathbf{in}\ \mathsf{bu}(F;\, e_2)) \\
\mathsf{bu}(F;\, \mathbf{if\ let}_\Delta \langle p \rangle = e_1\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3) &\coloneqq F(\mathbf{if\ let}_\Delta \langle p \rangle = \mathsf{bu}(F;\, e_1)\ \mathbf{then}\ \mathsf{bu}(F;\, e_2)\ \mathbf{else}\ \mathsf{bu}(F;\, e_3)) \\
\mathsf{bu}(F;\, e_1\ \mathbf{rewrite}\ \langle p \rangle \mapsto e_2) &\coloneqq F(\mathsf{bu}(F;\, e_1)\ \mathbf{rewrite}\ \langle p \rangle \mapsto \mathsf{bu}(F;\, e_2))
\end{aligned}
$$

$\boxed{\mathsf{bu}(F;\sigma)}$ *(Bottom-up Application on Substitutions)*

$$
\begin{aligned}
\mathsf{bu}(F;\, \cdot) &\coloneqq \cdot \\
\mathsf{bu}(F;\, \sigma, x \mapsto y) &\coloneqq \mathsf{bu}(F;\, \sigma),\ x \mapsto y \\
\mathsf{bu}(F;\, \sigma, x^m_\Delta \mapsto e) &\coloneqq \mathsf{bu}(F;\, \sigma),\ x^m_\Delta \mapsto \begin{cases} \mathsf{bu}(F;\, e) & \text{if } m = n, \\ e & \text{if } m > n. \end{cases}
\end{aligned}
$$

To ensure that F can be applied to any sub-expression, F must accept any expression of type ${\Gamma'}^n \vdash^n e' : A'$, where Γ' extends the context Γ of e, and A' is any level-n type. The output $F(e')$ must have the same type as e'.

As another example for rewriting strategies, we can define a top-most rewriting strategy $\mathsf{rewrite}_{\mathsf{TopMost}}$ that rewrites the top-most occurrence of a pattern p within an expression e_1 with another expression e_2:

$\boxed{\mathsf{rewrite}_{\mathsf{TopMost}}(e_1; p; e_2)}$ *(Top-most Rewriting)*

$$
\mathsf{rewrite}_{\mathsf{TopMost}}(e_1; p; e_2) = \begin{cases} e_2[\mathsf{id}_\Gamma, \sigma] & \text{if } A = B \text{ and } \mathsf{match}(p; e_1) = \sigma, \\ \mathsf{rewriteSubterms}(e_1; p; e_2) & \text{otherwise.} \end{cases}
$$

where $\mathsf{rewriteSubterms}(e_1; p; e_2)$ applies $\mathsf{rewrite}_{\mathsf{TopMost}}$ to immediate sub-expressions of e_1.

C Denotational Semantics Details

$\boxed{(\!|\, e \,|\!)_{\Gamma'}}$ *(Expression Interpretation)*

$$
\begin{aligned}
(\!|\, x_{\sigma_1} \,|\!)_{\Gamma'} &\coloneqq \lambda\rho.\ \rho(x)\ \mathsf{id}_{\Gamma'}\ (\!|\, \sigma_1 \,|\!)_{\Gamma'} \\
(\!|\, \mathbf{true} \,|\!)_{\Gamma'} &\coloneqq \lambda\rho.\ \mathsf{True} \\
(\!|\, \mathbf{false} \,|\!)_{\Gamma'} &\coloneqq \lambda\rho.\ \mathsf{False} \\
(\!|\, \mathbf{if}\ e_1\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3 \,|\!)_{\Gamma'} &\coloneqq \lambda\rho.\ \begin{cases} (\!|\, e_2 \,|\!)_{\Gamma'}\ \rho & \text{if } (\!|\, e_1 \,|\!)_{\Gamma'}\ \rho = \mathsf{True} \\ (\!|\, e_3 \,|\!)_{\Gamma'}\ \rho & \text{if } (\!|\, e_1 \,|\!)_{\Gamma'}\ \rho = \mathsf{False} \end{cases} \\
(\!|\, \lambda_\Delta x : A.\ e \,|\!)_{\Gamma'} &\coloneqq \lambda\rho.\ \lambda\sigma'\ d.\ (\!|\, e \,|\!)_{\Gamma'}\ (\rho[\sigma'] \cup \{x^n \mapsto d\}) \\
(\!|\, e_1\ e_2 \,|\!)_{\Gamma'} &\coloneqq \lambda\rho.\ (\!|\, e_1 \,|\!)_{\Gamma'}\ \rho\ \mathsf{id}_{\Gamma'}\ ((\!|\, e_2 \,|\!)_{\Gamma',\Delta}\ (\rho \cup \mathsf{id}_\Delta)) \\
(\!|\, \langle e \rangle \,|\!)_{\Gamma'} &\coloneqq \lambda\rho.\ e[\rho\restriction_{n+1}] \\
(\!|\, \mathbf{let}\ \langle \Delta.\ x \rangle = e_1\ \mathbf{in}\ e_2 \,|\!)_{\Gamma'} &\coloneqq \lambda\rho.\ \mathsf{let}\ e = (\!|\, e_1 \,|\!)_{\Gamma',\Delta}\ (\rho \cup \mathsf{id}_\Delta)\ \mathsf{in}\ (\!|\, e_2 \,|\!)_{\Gamma'}\ (\rho \cup (x \mapsto e)) \\
(\!|\, \mathbf{wrap}_\Delta e \,|\!)_{\Gamma'} &\coloneqq \lambda\rho.\ (\!|\, e \,|\!)_{\Gamma',\Delta}\ (\rho \cup \mathsf{id}_\Delta) \\
(\!|\, \mathbf{let\ wrap}_\Delta\ x : A = e_1\ \mathbf{in}\ e_2 \,|\!)_{\Gamma'} &\coloneqq \lambda\rho.\ \mathsf{let}\ d = (\!|\, e_1 \,|\!)_{\Gamma'}\ \rho\ \mathsf{in}\ (\!|\, e_2 \,|\!)_{\Gamma'}\ (\rho \cup \{x^n \mapsto d\})
\end{aligned}
$$

For $\lambda^{\bigcirc\triangleright}_{\mathrm{pat}}$ expressions:

$$
\begin{aligned}
(\!|\, \mathbf{if\ let}_\Delta \langle p \rangle = e_1\ \mathbf{then}\ e_2\ \mathbf{else}\ e_3 \,|\!)_{\Gamma'} &\coloneqq \lambda\rho.\ \begin{cases} (\!|\, e_2 \,|\!)_{\Gamma'}\ (\rho \cup \sigma) & \text{if } \mathsf{match}\big(p[\rho\restriction_{n+1}];\ (\!|\, e_1 \,|\!)_{\Gamma',\Delta}\ (\rho \cup \mathsf{id}_\Delta)\big) = \sigma, \\ (\!|\, e_3 \,|\!)_{\Gamma'}\ \rho & \text{otherwise.} \end{cases} \\
(\!|\, e_1\ \mathbf{rewrite}\ \langle p_1 \rangle \mapsto e_2 \,|\!)_{\Gamma'} &\coloneqq \lambda\rho.\ \mathsf{rewrite}_{\mathsf{BottomUp}}(p_1[\rho\restriction_{n+1}];\ (\!|\, e_2 \,|\!)_{\Gamma',\Pi}\ (\rho \cup \mathsf{id}_\Pi);\ (\!|\, e_1 \,|\!)_{\Gamma'}\ \rho)
\end{aligned}
$$

$\boxed{(\!|\,\sigma\,|\!)_{\Gamma'}}$ *(Substitution Interpretation)*

$$(\!|\,\sigma\,|\!)_{\Gamma'} \coloneqq \lambda\rho.\ \lambda(x : [\Delta' \vdash^m A] \in \Gamma).\ \begin{cases} \lambda\sigma'\ \rho'.\ (\!|\,e\,|\!)_{\Gamma'}\ (\rho[\sigma'] \cup \rho') & \text{if } m = n \text{ and } \sigma(x) = e, \\ e[\rho\restriction_m, \mathrm{id}_{\Delta'}] & \text{if } m > n \text{ and } \sigma(x) = e, \\ \rho(y) & \text{if } \sigma(x) = y, \end{cases}$$

References

Andreas Abel, Guillaume Allais, Aliya Hameer, Brigitte Pientka, Alberto Momigliano, Steven Schäfer, and Kathrin Stark. 2019. POPLMark reloaded: Mechanizing proofs by logical relations. *Journal of Functional Programming* 29 (Jan. 2019), e19. doi:10.1017/S0956796819000170

Emmanuel Suárez Acevedo and Stephanie Weirich. 2023. Making Logical Relations More Relatable (Proof Pearl). *arXiv preprint arXiv:2309.15724* (2023). doi:10.48550/arXiv.2309.15724

Thorsten Altenkirch, Martin Hofmann, and Thomas Streicher. 2005. Categorical reconstruction of a reduction free normalization proof. In *Category Theory and Computer Science*. Springer, Berlin, Germany, 182–199. doi:10.1007/3-540-60164-3_27

Kenichi Asai, Luminous Fennell, Peter Thiemann, and Yang Zhang. 2014. A type theoretic specification of partial evaluation. In *Proceedings of the 16th International Symposium on Principles and Practice of Declarative Programming*. 57–68. doi:10.1145/2643135.2643146

Eli Barzilay, Ryan Culpepper, and Matthew Flatt. 2011. Keeping it clean with syntax parameters. *Proc. Wksp. Scheme and Functional Programming* (2011).

G. M. Bierman and V. C. V. de Paiva. 2000. On an Intuitionistic Modal Logic. *Studia Logica* 65, 3 (Aug. 2000), 383–416. doi:10.1023/A:1005291931660

Mathieu Boespflug and Brigitte Pientka. 2011. Multi-level Contextual Type Theory. In *Proceedings Sixth International Workshop on Logical Frameworks and Meta-languages: Theory and Practice, LFMTP 2011, Nijmegen, The Netherlands, August 26, 2011 (EPTCS, Vol. 71)*, Herman Geuvers and Gopalan Nadathur (Eds.). 29–43. doi:10.4204/EPTCS.71.3

Cristiano Calcagno, Walid Taha, Liwen Huang, and Xavier Leroy. 2003. Implementing multi-stage languages using ASTs, gensym, and reflection. In *International Conference on Generative Programming and Component Engineering*. Springer, 57–76. doi:10.1007/978-3-540-39815-8_4

Tsung-Ju Chiang and Ningning Xie. 2025. *Multi-Stage Programming with Splice Variables (Artifact)*. doi:10.5281/zenodo.15719807

Tsung-Ju Chiang, Jeremy Yallop, Leo White, and Ningning Xie. 2024. Staged Compilation with Module Functors. *Proc. ACM Program. Lang.* 8, ICFP, Article 260 (Aug. 2024), 35 pages. doi:10.1145/3674649

William Clinger. 1991. Hygienic macros through explicit renaming. *ACM SIGPLAN Lisp Pointers* 4, 4 (1991), 25–28. doi:10.1145/1317265.1317269

Ranald Clouston. 2018. Fitch-style modal lambda calculi. In *Foundations of Software Science and Computation Structures: 21st International Conference, FOSSACS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14–20, 2018. Proceedings 21*. Springer, 258–275. doi:10.1007/978-3-319-89366-2_14

Rowan Davies. 1996. A temporal-logic approach to binding-time analysis. In *Proceedings 11th Annual IEEE Symposium on Logic in Computer Science*. IEEE, 184–195. doi:10.1109/LICS.1996.561317

Rowan Davies. 2017. A Temporal Logic Approach to Binding-Time Analysis. *J. ACM* 64, 1, Article 1 (March 2017), 45 pages. doi:10.1145/3011069

Marcelo Fiore and Chung-Kil Hur. 2010. Second-order equational logic. In *Computer Science Logic: 24th International Workshop, CSL 2010, 19th Annual Conference of the EACSL, Brno, Czech Republic, August 23-27, 2010. Proceedings 24*. Springer, 320–335. doi:10.1007/978-3-642-15205-4_26

Murdoch J. Gabbay and Aleksandar Nanevski. 2013. Denotation of contextual modal type theory (CMTT): Syntax and meta-programming. *Journal of Applied Logic* 11, 1 (2013), 1–29. doi:10.1016/j.jal.2012.07.002

Daniel Gratzer, G. A. Kavvos, Andreas Nuyts, and Lars Birkedal. 2020. Multimodal Dependent Type Theory. In *ACM Conferences*. Association for Computing Machinery, New York, NY, USA, 492–506. doi:10.46298/lmcs-17(3:11)2021

Jason Z. S. Hu and Brigitte Pientka. 2024. Layered modal type theory: Where meta-programming meets intensional analysis. In *European Symposium on Programming*. Springer, 52–82. doi:10.1007/978-3-031-57262-3_3

Jason Z. S. Hu and Brigitte Pientka. 2025. A Dependent Type Theory for Meta-programming with Intensional Analysis. *Proceedings of the ACM on Programming Languages* 9, POPL (2025), 416–445. doi:10.1145/3704851

Junyoung Jang, Samuel Gélineau, Stefan Monnier, and Brigitte Pientka. 2022. Mœbius: metaprogramming using contextual types: the stage where system f can pattern match on itself. *Proceedings of the ACM on Programming Languages* 6, POPL (Jan. 2022), 1–27. doi:10.1145/3498700

Manohar Jonnalagedda, Thierry Coppey, Sandro Stucki, Tiark Rompf, and Martin Odersky. 2014. Staged Parser Combinators for Efficient Data Processing. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming*

Systems Languages & Applications (Portland, Oregon, USA) *(OOPSLA '14)*. Association for Computing Machinery, New York, NY, USA, 637–653. doi:10.1145/2660193.2660241

G. A. Kavvos. 2017a. On the Semantics of Intensionality. In *Proceedings of the 20th International Conference on Foundations of Software Science and Computation Structures-Volume 10203*. 550–566. doi:10.1007/978-3-662-54458-7_32

G. A. Kavvos. 2017b. *On the Semantics of Intensionality and Intensional Recursion*. Ph. D. Dissertation. University of Oxford.

G. A. Kavvos. 2020. Dual-context calculi for modal logic. *Logical Methods in Computer Science* 16 (2020). doi:10.23638/LMCS-16(3:10)2020

G. A. Kavvos. 2024. Two-Dimensional Kripke Semantics I: Presheaves. *Schloss Dagstuhl – Leibniz-Zentrum für Informatik* (2024), 14:1–14:23. doi:10.4230/LIPIcs.FSCD.2024.14

G. A. Kavvos and Daniel Gratzer. 2023. UNDER LOCK AND KEY: A PROOF SYSTEM FOR A MULTIMODAL LOGIC. *Bulletin of Symbolic Logic* 29, 2 (June 2023), 264–293. doi:10.1017/bsl.2023.14

Oleg Kiselyov. 2014. The Design and Implementation of BER MetaOCaml. In *Functional and Logic Programming*, Michael Codish and Eijiro Sumii (Eds.). Springer International Publishing, Cham, 86–102. doi:10.1007/978-3-319-07151-0_6

Oleg Kiselyov, Aggelos Biboudis, Nick Palladinos, and Yannis Smaragdakis. 2017. Stream fusion, to completeness. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages*. 285–299. doi:10.1145/3093333.3009880

Eugene Kohlbecker, Daniel P. Friedman, Matthias Felleisen, and Bruce Duba. 1986. Hygienic macro expansion. In *Proceedings of the 1986 ACM Conference on LISP and Functional Programming*. 151–161. doi:10.1145/319838.319859

Kensuke Kojima and Atsushi Igarashi. 2011. Constructive linear-time temporal logic: Proof systems and Kripke semantics. *Information and Computation* 209, 12 (Dec. 2011), 1491–1503. doi:10.1016/j.ic.2010.09.008

John C. Mitchell and Eugenio Moggi. 1991. Kripke-style models for typed lambda calculus. *Annals of Pure and Applied Logic* 51, 1 (March 1991), 99–124. doi:10.1016/0168-0072(91)90067-V

Yuito Murase. 2017. Kripke-style contextual modal type theory. *Work-in-progress report at Logical Frameworks and Meta-Languages* (2017).

Yuito Murase, Yuichi Nishiwaki, and Atsushi Igarashi. 2023. Contextual Modal Type Theory with Polymorphic Contexts. In *Programming Languages and Systems*. Springer, Cham, Switzerland, 281–308. doi:10.1007/978-3-031-30044-8_11

Aleksandar Nanevski, Frank Pfenning, and Brigitte Pientka. 2008. Contextual modal type theory. *ACM Trans. Comput. Logic* 9, 3, Article 23 (June 2008), 49 pages. doi:10.1145/1352582.1352591

Lionel Parreaux, Antoine Voizard, Amir Shaikhha, and Christoph E. Koch. 2017. Unifying analytic and statically-typed quasiquotes. *Proceedings of the ACM on Programming Languages* 2, POPL (Dec. 2017), 1–33. doi:10.1145/3158101

Frank Pfenning and Rowan Davies. 2001. A judgmental reconstruction of modal logic. *Mathematical Structures in Computer Science* 11, 4 (Aug. 2001), 511–540. doi:10.1017/S0960129501003322

Tiark Rompf and Martin Odersky. 2010. Lightweight modular staging: a pragmatic approach to runtime code generation and compiled DSLs. In *Proceedings of the Ninth International Conference on Generative Programming and Component Engineering* (Eindhoven, The Netherlands) *(GPCE '10)*. Association for Computing Machinery, New York, NY, USA, 127–136. doi:10.1145/1868294.1868314

Tiark Rompf and Martin Odersky. 2012. Lightweight modular staging: a pragmatic approach to runtime code generation and compiled DSLs. *Commun. ACM* 55, 6 (June 2012), 121–130. doi:10.1145/2184319.2184345

Tim Sheard and Simon Peyton Jones. 2002. Template Meta-Programming for Haskell. In *Proceedings of the 2002 ACM SIGPLAN Workshop on Haskell* (Pittsburgh, Pennsylvania) *(Haskell '02)*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/581690.581691

Nicolas Stucki, Aggelos Biboudis, and Martin Odersky. 2018. A practical unification of multi-stage programming and macros. In *Proceedings of the 17th ACM SIGPLAN International Conference on Generative Programming: Concepts and Experiences*. 14–27. doi:10.1145/3278122.3278139

Nicolas Stucki, Jonathan Immanuel Brachthäuser, and Martin Odersky. 2021. Multi-stage programming with generative and analytical macros. In *ACM Conferences*. Association for Computing Machinery, New York, NY, USA, 110–122. doi:10.1145/3486609.3487203

Walid Taha and Tim Sheard. 2000. MetaML and multi-stage programming with explicit annotations. *Theoretical computer science* 248, 1-2 (2000), 211–242. doi:10.1016/S0304-3975(00)00053-0

The Agda Community. 2024. *Agda Standard Library*. https://github.com/agda/agda-stdlib

Philip Wadler, Wen Kokke, and Jeremy G. Siek. 2022. *Programming Language Foundations in Agda*.

Fei Wang, Daniel Zheng, James Decker, Xilun Wu, Grégory M Essertel, and Tiark Rompf. 2019. Demystifying differentiable programming: Shift/reset the penultimate backpropagator. *Proceedings of the ACM on Programming Languages* 3, ICFP (2019), 1–31. doi:10.1145/3341700

Jamie Willis, Nicolas Wu, and Matthew Pickering. 2020. Staged selective parser combinators. *Proc. ACM Program. Lang.* 4, ICFP (2020), 120:1–120:30. doi:10.1145/3409002

Ningning Xie, Matthew Pickering, Andres Löh, Nicolas Wu, Jeremy Yallop, and Meng Wang. 2022. Staging with class: a specification for typed template Haskell. *Proceedings of the ACM on Programming Languages* 6, POPL (Jan. 2022), 1–30. doi:10.1145/3498723

Ningning Xie, Leo White, Olivier Nicole, and Jeremy Yallop. 2023. MacoCaml: Staging Composable and Compilable Macros. *MacoCaml: Staging Composable and Compilable Macros (Artifact)* 7, ICFP (Aug. 2023), 604–648. doi:10.1145/3607851

Jeremy Yallop. 2017. Staged Generic Programming. *Proc. ACM Program. Lang.* 1, ICFP, Article 29 (Aug. 2017), 29 pages. doi:10.1145/3110273

Jeremy Yallop, Ningning Xie, and Neel Krishnaswami. 2023. flap: A Deterministic Parser with Fused Lexing. *Proc. ACM Program. Lang.* 7, PLDI, Article 155 (June 2023), 24 pages. doi:10.1145/3591269

Received 2025-02-27; accepted 2025-06-27