Lazy Linearity for a Core Functional Language

RODRIGO MESQUITA, Well-Typed LLP, United Kingdom BERNARDO TONINHO, Instituto Superior Técnico, University of Lisbon, Portugal and INESC-ID, Portugal

Traditionally, in linearly typed languages, consuming a linear resource is synonymous with its syntactic occurrence in the program. However, under the lens of non-strict evaluation, linearity can be further understood semantically, where a syntactic occurrence of a resource does not necessarily entail using that resource when the program is executed. While this distinction has been largely unexplored, it turns out to be inescapable in Haskell's optimising compiler, which heavily rewrites the source program in ways that break syntactic linearity but preserve the program's semantics. We introduce Linear Core, a novel system which accepts the lazy semantics of linearity statically and is suitable for lazy languages such as the Core intermediate language of the Glasgow Haskell Compiler. We prove that Linear Core is sound, guaranteeing linear resource usage, and that multiple optimising transformations preserve linearity in Linear Core while failing to do so in Core. We have implemented Linear Core as a compiler plugin to validate the system against linearity-heavy libraries, including linear-base.

CCS Concepts: • Theory of computation \rightarrow Linear logic; Type theory; • Software and its engineering \rightarrow Functional languages; Compilers.

Additional Key Words and Phrases: Linear Types, Lazy Evaluation, Haskell, GHC Core

ACM Reference Format:

Rodrigo Mesquita and Bernardo Toninho. 2026. Lazy Linearity for a Core Functional Language. *Proc. ACM Program. Lang.* 10, POPL, Article 69 (January 2026), 55 pages. https://doi.org/10.1145/3776711

1 Introduction

Linear types [Barber 1996; Girard 1987] increase the expressiveness of type systems by guaranteeing that certain resources are used exactly once. In programming languages with a linear type system, discarding so-called linear resources, or using them twice, is flagged as a type error. Linear types can be used to, for instance, statically enforce correct usage of file descriptors, that heap-allocated memory is freed exactly once s.t. leaks and double-frees become type errors, or guarantee deadlock freedom in channel-based communication protocols [Caires and Pfenning 2010], among other correctness properties [Bernardy et al. 2017a; Fu et al. 2020; Li et al. 2022]. The following program with a runtime error (allocated memory is freed twice) is accepted by a C-like type system. Conversely, under the lens of a linear type system where p is deemed a linear resource created by the call to malloc, the program is rejected:

let p = malloc(4) in free(p); free(p);

Despite their promise and extensive presence in research literature [Andreoli 1992; Cervesato et al. 2000; Wadler 1990], an effective design combining linear and non-linear typing is both challenging and necessary to bring the advantages of linear typing to mainstream languages. Consequently, few

Authors' Contact Information: Rodrigo Mesquita, Well-Typed LLP, London, United Kingdom, rodrigo@well-typed.com; Bernardo Toninho, Instituto Superior Técnico, University of Lisbon, Lisbon, Portugal and INESC-ID, Lisbon, Portugal, bernardo.toninho@tecnico.ulisboa.pt.



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

© 2026 Copyright held by the owner/author(s).

ACM 2475-1421/2026/1-ART69

https://doi.org/10.1145/3776711

general purpose programming languages have linear or substructural type systems. Among them are Idris 2 [Brady 2021], Rust [Matsakis and Klock 2014], whose ownership types build on linear types to guarantee memory safety without garbage collection, and, more recently, Haskell [Bernardy et al. 2017b], a pure, functional, and *lazy* general purpose language.

Linearity in Haskell stands out from linearity in other languages because linear types permeate Haskell down to Core, the intermediate language into which Haskell is translated by the Glasgow Haskell Compiler (GHC). Core is a minimal, explicitly typed, pure lazy functional language with both linear and unrestricted types, to which multiple Core-to-Core optimising transformations are applied during compilation. Notably, Core is typechecked after each transformation. This serves as a sanity check to the correction of the implementation of such transformations, since unsound ones are likely to introduce ill-typed expressions.

Just as Core's type system [Sulzmann et al. 2007] provides a degree of validation to the translation from Haskell, dubbed *desugaring*, and the subsequent optimising transformations, a linearly typed Core would guarantee that linear resource usage in the source language is not violated by desugaring and the compiler optimisation passes. Moreover, linearity information in Core can inform Core-to-Core optimising transformations [Bernardy et al. 2017b; Peyton Jones and Partain 1996; Peyton Jones and Santos 1997], such as inlining and β -reduction, to produce more efficient programs. However, despite the advantages of having linear types in Core, the status quo is that linearity is effectively ignored in Core, being only checked in the source Haskell code. The reason is that, after desugaring, the Core-to-Core optimising transformations eliminate and rearrange most of the syntactic constructs through which linearity checking is performed – often resulting in programs that are syntactically very different from the original, especially due to the flexibility provided by laziness with regard to the optimisations a compiler may perform.

While compiler optimisations aim to preserve the semantics of programs and therefore a program's resource consumption, a naive syntactic analysis of the optimised programs as performed by the current Core type checker often fails to recognize resulting programs as linear. As a trivial example, let x be a linear resource in the next two expressions, where the latter results from inlining y without eliminating the binder – a common optimising transformation in GHC. Even though the second expression no longer appears linear (as there are now two occurrences of x), it is indeed linear because the let-bound expression is never evaluated under non-strict evaluation and x is still consumed exactly once (when it is freed in the let body):

let
$$y = \text{free } x \text{ in } y \implies_{Inlining} \text{let } y = \text{free } x \text{ in free } x$$

The Core optimising transformations expose a fundamental limitation of Core's linear type system: it does not account for the non-strict evaluation model of Core and, thus, a whole class of programs that are linear under the lens of non-strict evaluation are rejected by Core's current checker. In this work, we address this limitation by exploring how syntactic linearity breaks down in the presence of non-strictness and design a linear type system for Core which accepts the vast majority of programs and optimising transformations that, while correct under non-strict evaluation, are rejected by traditional linear typing disciplines. Additionally, we show that our system is suitable for the intermediate language of an optimising compiler by implementing it as a Glasgow Haskell Compiler plugin. Our contributions are as follows:

We illustrate the lazy semantics of linearity, in contrast to the traditional syntactic understanding of linearity, in Haskell, by example (Section 2). A precise semantic definition of linearity, closely following that of Linear Haskell [Bernardy et al. 2017a], is given by our instrumented operational semantics under which a program using a linear resource non-linearly gets stuck (Section 4).

- We introduce Linear Core (Section 3), a lazy linear language with the key features of Core (except for type equality coercions), whose type system is designed to accept the lazy semantics of linearity. Our system achieves this via a novel combination of features: explicit tracking of captured linear resources through *usage environments* (Section 3.3); a distinct treatment of case scrutinees in weak-head normal form (Section 3.6.1); and a notion of irrelevant and tagged resources (Section 3.6.2), which allow for fine-grained control over variable usage. To the best of our knowledge, this is the first type system to accept programs which depend on non-strict evaluation to be understood as linear.
- We prove that our type system is sound, guarantees linear resource usage, and prove that multiple optimising transformations, which are rejected in Core, are validated by Linear Core (Section 4).
- We discuss our implementation of Linear Core as a GHC plugin which checks linearity in all intermediate Core programs produced during the compilation process, showing it accepts the programs resulting from transformations in libraries such as linear-base (Section 5).

While GHC's Core language is our main focus, we note that the ideas explored in this work can potentially be applied to any language with non-strict features, even if the language itself is strict. We assume some familiarity with Haskell and linear types for functional languages, but try to explain our work and motivations without assuming much knowledge of GHC and Core.

Additional proofs and definitions are given in Appendix.

2 Linearity, Semantically

Linear type systems guarantee that linear resources are consumed exactly once. Thus, whether a program is accepted as linear intrinsically depends on the definition of consuming a resource. Traditionally, consuming a resource has been equated with its syntactic occurrence along a control-flow path in the program. However, as this section aims to make clear, in a non-strict context, a syntactic definition of linearity is overly conservative. For instance, in the following Haskell-like program, f is a linear function (a function that must use its argument exactly once), but there are two syntactic occurrences of handle: one is in a let-bound computation that closes this handle, the other is in the then branch of a conditional. If the program is evaluated lazily, the handle will, dynamically, either be used exactly once (i.e, closed) in the else branch, or be used exactly once (i.e., returned) in the then branch, since the let-bound variable is not needed in that branch and so the expression which closes the handle is never evaluated.

```
f: Handle - Handle
f handle =
  let closed = close handle in
  if <cond>
    then handle
  else closed
```

Intuitively, a computation that depends on a linear resource to produce some result consumes that resource iff the result is fully consumed; in contrast, a computation that depends on a linear resource, but is never run, will not consume that resource. We argue that a linear type system for a language with non-strict evaluation semantics should accept the above program, unlike a linear type system for the same program if it were to be evaluated eagerly, where the let-bound expression would always be evaluated. Given this observation, we turn our focus to linearity under lazy evaluation specifically since the distinction between semantically and syntactically consuming a resource is only exposed under non-strict semantics, and because GHC Core is lazy.

2.1 Semantic Linearity by Example

This section aims to develop an intuition for semantic linearity through simple Core programs that can be understood as linear but are rejected by Core's linear type system (or, in fact, by all linear type systems we are aware of prior to this work). In the examples, a light green background highlights programs that are syntactically linear and accepted by Core's linear type system. A light yellow or light orange background mark programs that are semantically linear, but not seen as linear by Core's type system. Our system accepts all light yellow programs, but not the light orange examples. A light red background indicates that the program violates linearity semantically, i.e. the program effectively discards or duplicates linear resources.

Let bindings. The example at the start of Section 2 highlighted the subtlety of linearity in a non-strict language with a single unused let binding. In general, a let-bound expression that captures a linear variable can be conditionally needed at runtime, depending on the branch taken. Both optimising transformations (float-out) and programmers used to non-strict evaluation are keen to write programs with bindings that are selectively used in case alternatives, such as in the following program. Note that we write \multimap for linear functions and \multimap for non-linear functions:

```
f :: (a -\infty a) -> Bool -> a -\infty a
f use bool x = let y = use x in
case bool of
True -> x
False -> y
```

This example returns x in one branch and the let-bound y in the other. Semantically, as opposed to syntactically, this program is linear because the linear resource x is used exactly once regardless of which branch is taken: either directly in the **True** branch, or *via* the use of the let-bound y in the **False** branch.

That said, a let-bound y that captures a linear resource x must still be used exactly once, regardless of whether evaluation inlines the let (\dot{a} la call-by-name) or creates a thunk (\dot{a} la call-by-need):

```
f1 :: (a \multimap b) \multimap a \multimap (b, b)
f1 use x = let y = use x in (y, y)
```

If y is inlined, linearity is trivially violated. On the other hand, if y becomes a thunk, use x will only be evaluated once regardless of how many times y is used. Even so, this program must be rejected under call-by-need. Intuitively, linearity is violated if the computation leaks the resource or parts of it. For instance, if use = id, the program must clearly be rejected.

Lastly, consider an ill-typed program which defines two let bindings z and y, where z uses y which in turn uses the linear resource x:

```
f2 :: (a \multimap a) \multimap a \multimap ()
f2 use x = let y = use x in let z = use y in ()
```

Even though y occurs in z, x is never consumed because z is discarded. This example highlights that naively checking for syntactic occurrences of let-bound variables such as y is also inadequate to track resource usage. A better mental model is that using y implies using x, and thus using y implies using y too. If neither y nor y are used, then y is effectively discarded.

In essence, an unused let binding doesn't consume any resources, and a let binding used exactly once consumes the resources it captures exactly once. Let binders that depend on linear resources must be used *at most once* in mutual exclusion with the linear resources themselves (in a sense,

let-bound variables are affine in the let body). We discuss how to encode this principle between let bindings and their dependencies using so called *usage environments*, in Section 3.3.

Recursive let bindings. Recursive let-bindings behave similarly to non-recursive lets as far as their usage in the let continuation is concerned. The challenge lies in understanding linearity in the definitions of the recursive binders and determining which resources are consumed when one of the binders in the recursive group is used. Consider the following program, which calls a recursive let-bound function defined in terms of the linear resource x and itself:

```
f3 :: Bool -> a -o a
f3 bool x = let go b = case b of

True -> x

False -> go (not b)

in go bool
```

Function f3 is semantically linear because, iff it is consumed exactly once, then x is consumed exactly once. We can see this by case analysis on go's argument: when bool is **True**, we'll use the resource x; when bool is **False**, we recurse by calling go on **True**, which in turn will use the resource x. In go's body, x is used directly in one branch and indirectly in the other (by recursing, which we know will result in using x linearly).

It so happens that go will terminate on any input, and will always consume x. However, termination is not a requirement for a binding to use a resource linearly, and we could have a similar example in which go might never terminate but still uses x linearly if evaluated:

```
f4 :: Bool -> a -∞ a

f4 bool x = let go b = case b of

True -> x

False -> go b

in go bool
```

The key to linearity in the presence of non-termination is Linear Haskell's definition of a linear function: if a linear function application $(f \ u)$ is consumed exactly once, then the argument (u) is consumed exactly once. If f u doesn't terminate, it is never consumed, thus the claim holds vacuously. If go doesn't terminate, we aren't able to compute (or consume) the result of f4, so we do not promise anything about x being consumed (f4's linearity holds trivially). If it did terminate, it would consume x exactly once (e.g. if go was applied to **True**).

Naturally, a linear type system must statically reject any recursive group that could lead to non-linear resource usage. For instance, if the linear resource x is used in one case alternative and recursion is used more than once in another, x may end up being consumed more than once:

```
f5 :: Bool -> Bool -0 Bool

f5 bool x =

let go b

= case b of

True -> x

False -> go (not b) && go True

in go bool
```

To determine whether the usage of a binding from a recursive group is linear, we must check the superset of resources that may end up being consumed in a single use of any of the bindings. For

a strongly connected recursive group, this superset is the same for all bindings by definition. It follows from the resources used being the same for all bindings that at most a single variable from the recursive group can be used in the continuation. The treatment of recursive let bindings in our system is discussed in Section 3.5.

Case expressions. We have observed that more programs can be accepted as linear by delaying resource consumption in (recursive) let bindings until they are evaluated. On the other hand, case expressions drive evaluation – and, therefore, do consume resources.

To understand exactly how, we turn to the definition of consuming a resource from Linear Haskell [Bernardy et al. 2017b]: to consume a value of atomic base type (such as Int or Ptr) exactly once, just evaluate it; to consume a function value exactly once, apply it to one argument, and consume its result exactly once; to consume a value of an algebraic datatype exactly once, pattern-match on it, and consume all its linear components exactly once. Thus, we can consume a linear resource by fully evaluating it, and evaluation happens with case expressions.

In Core, cases are of the form case e_s of z { $\overline{\rho_i} \to \overline{e_i}$ }, where e_s is the case *scrutinee*, z is the case *binder*, and $\overline{\rho_i} \to \overline{e_i}$ are the case *alternatives*, composed of a pattern ρ_i and of the corresponding expression e_i . The case scrutinee is always evaluated to Weak Head Normal Form (WHNF) and the case binder is an alias to the result of evaluating the scrutinee to WHNF. Moreover, in Core, the alternative patterns are always exhaustive. The first example program constructs a tuple from linear resources, matches on it, then uses both linearly-bound variables from the tuple pattern. This is well-typed in Linear Haskell:

```
f6 :: a \multimap b \multimap (a \multimap b \multimap c) \multimap c
f6 x y use = case (x,y) of z { (a,b) -> use a b }
```

Perhaps surprisingly, a similar program which discards the pattern variables, and instead uses the resources from the scrutinee again, is still linear, despite not being accepted by Linear Haskell:

```
f7 :: a \multimap b \multimap (a \multimap b \multimap c) \multimap c
f7 x y use = case (x,y) of z { (a,b) -> use x y }
```

f7 can be seen as linear by realizing that the tuple being scrutinized is already in WHNF, so evaluating it again does not consume either x nor y. Even if the tuple paired two arbitrary expressions which captured x and y, because a and b are not forced in the continuation, x and y would remain unused. On the other hand, as seen in the next example, if a were used in the case body, then x is consumed via a and must not be used again. Thus, f8 must be rejected:

```
f8 :: a \multimap b \multimap (a \multimap a \multimap c) \multimap c
f8 x y use = case (x,y) of z { (a,b) \multimap use a x }
```

This idea that x and a are mutually exclusive is akin to how let bindings are mutually exclusive with the set of resources they capture, in the let-continuation. Forcing a pattern variable evaluates the expression for the corresponding constructor field, and all linear variables captured by that expression may end up being used. A third option, semantically linear, but rejected by Linear Haskell, is to use the case binder z instead of a, b or x, y:

```
f9 :: a \multimap b \multimap (a \multimap b \multimap c) \rightarrow c
f9 x y use = case (x,y) of z { (a,b) -> uncurry use z }
```

Again, z is mutually exclusive with (a, b) and with (x, y), but at least one of the three must be used to ensure the linear resources are consumed. Essentially, in this example, using a entails using the resource x, b the resource y, and the case binder z entails using both a and b.

Conversely, if the scrutinee is an expression that's not in WHNF, evaluation to WHNF may indeed consume some or all of the linear resources captured by said scrutinee. Consider the following non-WHNF scrutinee example, f10, which is accepted by our system:

```
f10 :: a \multimap b \multimap (a \multimap b \multimap (c,d)) \multimap (c,d)
f10 x y use = case use x y of z { (a,b) \multimap z }
```

In contrast to a scrutinee in WHNF, the resources captured by a non-WHNF scrutinee (here, x, y), conservatively, can no longer be used in the case alternatives, since reducing use x y to WHNF may consume (parts of) x and/or y. Consequently, either the case binder z or the linear pattern variables a, b must be used exactly once, since the linear resources in the scrutinee are only fully consumed if the result of evaluating the scrutinee is also fully consumed . For instance, take use in f10 to be the pairing function (,): it is not sufficient for use x y to be evaluated to WHNF to consume x and y. If all the resources were considered to be fully consumed after evaluating the scrutinee to WHNF, the pattern variables could simply be ignored and thus linear resources potentially discarded. In practice, if the scrutinee is not in WHNF then either the case binder or the linear components of the pattern must be consumed.

Alternatively, if a constructor without any linear components is matched, all linear resources used in the scrutinee must already have been fully consumed in that branch since the result of evaluating it can be consumed unrestrictedly (by definition, a constructor application to unrestricted arguments is also unrestricted). Hence, in the case-branch for a constructor without linear fields, the case binder can also be used unrestrictedly. For example, the program in Figure 1 is semantically linear. A second example of an unrestricted pattern, where K2 has no fields and K1 has one linear

```
f11 :: () \rightarrow () f12 :: a \rightarrow a f12 x = case x of z { () -> z <> z } f12 x = case K1 x of z { K2 -> x; K1 a -> x }
```

Fig. 1. No Linear Fields

Fig. 2. Absurd Branches

field, is shown in Figure 2. The use of x in the K2 branch is valid despite x also occurring in the scrutinee because this branch is never taken. In fact, any arbitrary resource could be used in absurd branches and be understood as linear (despite not being seen as such by our system).

While many of these examples may seem artificial due to the pattern match on a known constructor, we note that in between the transformations programs undergo in an optimising compiler, many such programs naturally occur (e.g., if the definition of a function is inlined in the scrutinee).

As for *default* case alternatives, also known as *wildcards* (written _): matching a wildcard doesn't provide any linearity information, but the scrutinee is still evaluated to WHNF. When matching a wildcard, if the scrutinee is already in WHNF, programs that either use the resources from the scrutinee directly, or the case binder in that alternative, can be seen as linear. If the scrutinee is not in WHNF, we *must* use the case binder, as it's the only way to linearly consume the result of evaluating the scrutinee.

3 A Type System for Semantic Linearity in Core

In this section, we develop a pure linear calculus λ_{Δ}^{π} , dubbed *Linear Core*, that combines the linearity-in-the-arrow and multiplicity polymorphism introduced by Linear Haskell [Bernardy et al. 2017b]

with all the key features of GHC Core: algebraic datatypes, case expressions, and recursive lazy let bindings; except for equality coercions, which we discuss as future work in Section 6.1.

Linear Core makes precise the various insights discussed in the previous section in a linear type system that we prove to be sound and guarantee linear resource usage at runtime. This result is obtained (see Section 4) via a soundness argument on a call-by-need big-step operational semantics that becomes stuck when linear variables in the runtime heap are used more than once, a technical approach that is natural in the literature [Bernardy and Spiwack 2021; Pierce 2004]. We show this semantics to be bisimilar to a Launchbury-style [Launchbury 1993] natural semantics. Crucially, Linear Core typing accepts all the semantically linear examples (highlighted with light yellow) from Section 2.1, which Core currently rejects. Additionally, we prove that multiple optimising Core-to-Core transformations preserve linearity in Linear Core, while violating it under Core's current type system. Despite the focus on GHC Core, the fundamental ideas for typing linearity in our call-by-need calculus may be applied to other languages which combine non-strict features or semantics with linearity.

The key ideas that make Linear Core work are:

- (1) Usage environments (Section 3.3)
- (2) Distinct treatment of case scrutinees in WHNF (Section 3.6.1)
- (3) Irrelevant and tagged resources (Sections 3.6.2 and 3.6.3)
- (1) Encodes the idea that lazy bindings do not consume resources upon definition, but rather when the bound variables themselves are consumed. Consuming a variable with some usage environment Δ equates to consuming the variables contained in Δ . (2) Case expressions in WHNF may capture ambient linear resources, but these resources can be safely used in branches since the case will not further evaluate its scrutinee. On the other hand, scrutinizing a non-WHNF expression must be treated more conservatively. (3) Irrelevant resources make existing bindings unusable while forcing the use of others. Namely, linear resources occurring in non-WHNF case scrutinees can no longer be used in the alternatives, while the pattern variables (or the case binder) must necessarily be used. Tagged resources follow the observation that some variables must be used jointly or not at all.

The reader may wonder whether the complexities of our system hinted above are a better alternative to simply specializing the optimisation passes to produce code that passes a more naive linearity check. This avenue was already explored during the implementation of Linear Haskell in GHC, exhibiting the prohibitive cost of too often restricting optimisations and producing less efficient code. Our work adheres to the principle that the type system bends to the optimisations, not the other way around¹.

We present Linear Core's syntax and type system incrementally, starting with the judgements and base linear calculi rules (Section 3.2). Then, usage environments, the rule for Δ -bound variables, and rules for (recursive) let bindings (Section 3.3). Finally, we introduce the rules to type case expressions and case alternatives, along with the key insights to do so, namely discriminating by the WHNF-ness of the scrutinee, irrelevant resources, and tagging (Section 3.6).

3.1 Language Syntax

The complete syntax of Linear Core is given in Figure 3. The types of Linear Core are algebraic datatypes, function types, and multiplicity schemes to support multiplicity polymorphism: datatypes $(T \ \overline{p})$ are parametrised by multiplicities, function types $(\varphi \rightarrow_{\pi} \sigma)$ are also annotated with a multiplicity, which can be 1, ω (read many), or a multiplicity variable p introduced by a multiplicity universal scheme $(\forall p.\ \varphi)$.

 $^{^1} https://gitlab.haskell.org/ghc/ghc/-/blob/c85c845dc5ad539bf28f1b8c5c1dbb349e3f3d25/compiler/GHC/Core/Lint.hs\#L3225-3235$

Terms are variables x, y, z, data constructors K, multiplicity abstractions Λp . e and applications e π , term abstractions $\lambda x:_{\pi} \tau$. e and applications e e', where lambda binders are annotated with a multiplicity π and a type σ . Then, there are non-recursive let bindings let $x:_{\Delta}\sigma = e$ in e', recursive let bindings let $\operatorname{rec} \overline{x:_{\Delta}\sigma} = e$ in e', where the overline denotes a set of distinct bindings $x_1:_{\Delta}\sigma_1 \dots x_n:_{\Delta}\sigma_n$ and associated expressions $e_1 \dots e_n$, and case expressions case e of $z:_{\Delta}\sigma$ $\{\overline{\rho} \Rightarrow e'\}$, where z is the case binder and the overline denotes a set of distinct patterns $\rho_1 \dots \rho_n$ and corresponding right hand sides $e'_1 \dots e'_n$. We often omit the case binder z when it is unused. Notably, (recursive) let-bound binders and case-bound binders are annotated with a so-called usage environment Δ – a fundamental construct for typing semantic linearity in the presence of laziness that we present in Section 3.3. Case patterns ρ can be either the default/wildcard pattern _, which matches any expression, or a constructor K and a set of variables that bind its arguments, where each field of the constructor has an associated multiplicity denoting which pattern-bound variables must be consumed linearly or unrestrictedly. Additionally, the set of patterns in a case expression is assumed to be exhaustive, i.e. there is always at least one pattern which matches the scrutinized expression.

Datatype declarations data $T \, \overline{p}$ where $K : \overline{\sigma \to_{\pi}} \, T \, \overline{p}$ have the name of the type being declared T parametrized over multiplicity variables \overline{p} , and a set of the data constructors K with signatures indicating the type and multiplicity of the constructor arguments. Note that a linear resource is used many times when a constructor with an unrestricted field is applied to it, since, dually, pattern matching on the same constructor with an unrestricted field allows it to be used unrestrictedly.

```
Types
  \tau, \sigma ::= T \overline{p}
                                  Datatype
                  \tau \rightarrow_{\pi} \sigma Function with multiplicity
                  ∀p. τ
                                  Multiplicity universal scheme
Terms
     ::= x, y, z \mid K
  е
                                                      Variables and data constructors
        | \Lambda p. e | e \pi
                                                      Multiplicity abstraction/application
         | \lambda x:_{\pi} \tau. e | e e'
                                                      Term abstraction/application
             let x:_{\Lambda} \sigma = e in e'
                                                      Let with usage environment
               let rec \overline{x:_{\Delta}\sigma = e} in e'
                                                      Recursive Let with usage environment
               case e of z:_{\Delta} \tau \{ \overline{\rho \Rightarrow e'} \}
     := K \overline{x} :_{\pi} \tau |_{-}
                                                      Pattern and wildcard
Environments
  \Gamma ::= \cdot \mid \Gamma, x:_{\omega} \tau \mid \Gamma, K:\tau \mid \Gamma, p \mid \Gamma, z:_{\Delta} \tau Unrestricted typing environment
      := \cdot | \Delta, x_{\pi} \tau | \Delta, [x_{\pi} \tau]
                                                                 Linear typing environment
Multiplicities
                                   Declarations
  \pi ::= 1 \mid \omega \mid p
                                  decl ::= \mathbf{data} \ T \ \overline{p} \ \mathbf{where} \ K : \overline{\tau \to_{\pi}} \ T \ \overline{p}
```

Fig. 3. Linear Core Syntax

3.2 Base Typing System

Linear Core (λ_{Δ}^{π}) is a linear lambda calculus akin to Linear Haskell's λ_{Δ}^{q} in that both have multiplicity polymorphism, (recursive) let bindings, case expressions, and algebraic data types. We use the π superscript to stand for multiplicity and the Δ subscript to emphasize usage environments, a key feature of our work. λ_{Δ}^{π} diverges from λ_{Δ}^{q} primarily when typing lets, case expressions, and

alternatives. The core rules of the calculus for abstraction and application are similar to that of λ^q_{\rightarrow} and mostly standard. We note that we handle multiplicity polymorphism differently from Linear Haskell by effectively ignoring the multiplicity semiring and treating all multiplicity polymorphic variables parametrically, for the sake of simplicity. Such variables cannot be duplicated in our system. The full type system is given in Figure 4, with auxiliary judgements given in Figure 5.

$$\begin{array}{c} \boxed{\Gamma;\Delta \vdash e : \tau} \\ \\ \frac{\Gamma, \rho; \Delta \vdash e : \tau}{\Gamma; \Delta \vdash \Lambda \rho, e : \forall \rho, \tau} \quad (\Delta I) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \Delta \vdash \Lambda \rho, e : \forall \rho, \tau} \quad (\Delta I_1) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \Delta \vdash \lambda x_{:1} \tau, e : \tau} \quad (\Delta I_2) \\ \\ \frac{\Gamma; \Delta \vdash \lambda x_{:1} \tau, e : \tau}{\Gamma; \Delta \vdash \lambda x_{:2} \tau, e : \tau} \quad (\Delta I_1) \\ \\ \frac{\Delta = \Delta'}{\Gamma; \Delta \vdash \lambda x_{:1} \tau, e : \tau} \quad (\Delta I_2) \\ \\ \frac{\Delta = \Delta'}{\Gamma; \Delta \vdash \lambda x_{:1} \tau, e : \tau} \quad (\Delta I_2) \\ \\ \frac{\Gamma; \Delta \vdash \lambda x_{:2} \tau, e : \tau}{\Gamma; \Delta \vdash \lambda x_{:2} \tau} \quad (\Delta I_3) \\ \\ \frac{\rho \in \Gamma}{\Gamma; x_{:2} \tau \vdash x : \tau} \quad (\Delta I_4) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, x_{:1} \tau \vdash x : \tau} \quad (\Delta I_2) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda, \tau} \quad (\Delta I_3) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda, \tau} \quad (\Delta I_4) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda, \tau} \quad (\Delta I_4) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_2) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_3) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_4) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_5) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_5) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_5) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_5) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_5) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_7; \Delta, \lambda' \vdash e' : \tau) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_7; \Delta, \lambda' \vdash e' : \tau) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_7; \Delta, \lambda' \vdash e' : \tau) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_7; \Delta, \lambda' \vdash e' : \tau) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_7; \Delta, \lambda' \vdash e' : \tau) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_7; \Delta, \lambda' \vdash e' : \tau) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_7; \Delta, \lambda' \vdash e' : \tau) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_7; \Delta, \lambda' \vdash e' : \tau) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda' \vdash e' : \tau} \quad (\Delta I_7; \Delta, \lambda' \vdash e' : \tau) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda' \vdash e : \tau} \quad (\Delta I_7; \Delta, \lambda' \vdash e' : \tau) \\ \\ \frac{\Gamma; \Delta \vdash e : \tau}{\Gamma; \lambda, \lambda' \vdash e : \tau} \quad (\Delta I_7; \Delta, \lambda' \vdash e' : \tau) \\ \\ \frac{\Gamma; \lambda \vdash e : \tau}{\Gamma; \lambda \vdash e : \tau} \quad (\Delta I_7; \Delta, \lambda' \vdash e' : \tau) \\ \\ \frac{\Gamma; \lambda \vdash e : \tau}{\Gamma; \lambda \vdash e : \tau} \quad (\Delta I_7; \Delta, \lambda' \vdash e' : \tau) \\ \\ \frac{\Gamma; \lambda \vdash e : \tau}{\Gamma; \lambda \vdash e : \tau} \quad (\Delta I_7; \Delta, \lambda' \vdash e' : \tau) \\ \\ \frac{\Gamma; \lambda \vdash e : \tau}{\Gamma; \lambda \vdash e : \tau} \quad (\Delta I_7; \Delta, \lambda' \vdash e' : \tau) \\ \\ \frac{\Gamma; \lambda \vdash e : \tau}{\Gamma; \lambda \vdash e : \tau} \quad (\Delta I_7; \Delta, \lambda$$

Fig. 4. Linear Core Type System

$$\frac{\Gamma \vdash_{mult} \pi}{\Gamma \vdash_{\Gamma} \Gamma} (1) \qquad \frac{\Gamma \vdash_{mult} \pi}{\Gamma \vdash_{\Gamma} \Gamma} (1) \qquad \frac{\Gamma \vdash_{\Gamma} \Gamma}{\Gamma \vdash_{\Gamma} \Gamma} (1) \qquad \frac{\Gamma \vdash_{\Gamma} \Gamma}{\Gamma} (1)$$

Fig. 5. Linear Core Auxiliary Judgements

The main judgement is written Γ ; $\Delta \vdash e : \tau$ to denote that expression e has type τ under the unrestricted environment Γ and linear environment Δ . Variables in Γ can be freely discarded (weakened) and duplicated (contracted), while resources in Δ must be used exactly once. Despite not having explicit weakening and contraction rules in our system, they are available as admissible rules for Γ (but not for Δ), since, equivalently [Dyckhoff 1992], resources from Γ are duplicated for sub-derivations and may unrestrictedly exist in the variable rules.

Occurrences of unrestricted variables from Γ are well-typed as long as the linear environment is empty, while occurrences of linear (or multiplicity polymorphic) variables are only well-typed when the variable being typed is the only resource available in the linear context (rules Var_{ω} , Var_{1} and Var_{p} , respectively). In the Var_{1} and Var_{p} cases, the linear context must contain exactly the appropriate variable binding, whereas the unrestricted context may contain arbitrary variables. Variables in contexts are annotated with their type and multiplicity, so $x:_{\pi}\tau$ is a variable named x of type σ and multiplicity π .

Linear functions are introduced via the function type $(\sigma \to_{\pi} \varphi)$ with $\pi = 1$, i.e. a function of type $\sigma \to_{1} \varphi$ (or, equivalently, $\sigma \to \varphi$) introduces a linear resource of type σ in the linear environment Δ to then type an expression of type φ . Unrestricted functions are introduced via the function type $(\sigma \to_{\pi} \varphi)$ with $\pi = \omega$, and the λ -bound variable is introduced in Γ (rules λI_{1} and ΛI). Function application is standard, with linear function application splitting the linear context in two disjoint parts, one used to type the argument and the other the function. Arguments of unrestricted functions must also be unrestricted, i.e. no linear variables can be used to type them. Typing linear and unrestricted function application separately is less general than typing applications of functions of any multiplicity π by scaling (per the multiplicity semiring) the multiplicities of the resources used to type the argument by π , however, our objective of typing semantic linearity does not benefit much from doing so and therefore we opt for a simpler design.

Multiplicity abstractions (ΛI) introduce a multiplicity variable p, and construct expressions of type $\forall p$. σ , i.e. a type universally quantified over a multiplicity variable p. In the body of the abstraction, function types and datatype fields annotated with the multiplicity variable p are typed parametrically in their multiplicity, since p can be instantiated at both ω and 1. A multiplicity application (ΛE) instantiates a multiplicity-polymorphic type $\forall p$. σ at a particular (argument) multiplicity π , resulting in an expression of type σ where occurrences of p are substituted by π , i.e. $\sigma[\pi/p]$. The rule additionally requires that π be well-formed in order for the expression to be well-typed, using the judgement $\Gamma \vdash_{mult} \pi$, where well-formedness is given by π either being 1, ω , or an in-scope multiplicity variable in Γ .

3.3 Usage environments

A usage environment Δ encodes the idea that lazy variable bindings do not consume resources upon definition, but rather when the bindings themselves are consumed. Specifically, we annotate so-called Δ -bound variables with their usage environment Δ to denote that consuming such variables must be equated with consuming the resources in the annotated typing environment. Such variables are introduced in the unrestricted environment by multiple constructs such as (recursive) let binders, case binders, and case pattern variables. For example, the let-bound variable u in the following program is annotated with a usage environment $\{x:_1\sigma, y:_1\sigma\}$ tracking the resources that are used in its body. Thus, in the expression e, consuming u will equate to consuming the linear resources x and y:

$$f = \lambda x_{1} \tau_{1} \tau_{1} \lambda y_{1} \tau_{1}$$
. let $u_{\{x_{1} \tau, y_{1} \sigma\}} = (x, y)$ in e

Furthermore, usage environments guarantee that using a Δ -bound variable is mutually exclusive with directly using the resources it is annotated with – using the Δ -bound variable consumes all linear resources listed in its usage environment, meaning they are no longer available for direct usage. Dually, using the linear resources directly means they are no longer available to consume through the usage environment of the Δ -bound variable. The Δ -bound variable can be left unused since it is unrestricted, regardless of the variables it is annotated with. In the example above, our type system ensures that *either* x and y are consumed or y is consumed in y, since y has usage environment y in y in y in y.

 Δ -bound variables. A Δ -bound variable u is a variable annotated with some usage environment Δ' . For conciseness, we use the meta-variable Δ to refer to the usage environment of a Δ -bound variable. Crucially, for any Δ -bound variable u: (1) using u is equivalent to using all the linear resources in Δ ; (2) using u is mutually exclusive with using the Δ resources it depends on; (3) u can be safely discarded as long as the resources in Δ are consumed. Fortunately, since linear resources must be linearly split across sub-derivations, (2) follows from (1) since consuming the linear resources in Δ to type u makes them unavailable in the context of any other derivation. Similarly, (3) also follows from (1), because if the linear resources aren't consumed by a use of u, they must be consumed in some other derivation (or otherwise the overall expression is ill-typed). These observations lead to the typing rule for Δ -bound variables:

$$\frac{\Delta = \Delta'}{\Gamma, x:_{\Lambda} \sigma; \Delta' \vdash x : \tau} (VAR_{\Delta})$$

The rule states that an occurrence of a Δ -bound variable is well-typed if the linear environment is made up exactly of the resources in the usage environment of that variable, denoted by the $\Delta = \Delta'$ premise of the rule. Δ -bound variables are unrestricted and can be discarded and duplicated, despite multiple occurrences of the variable in an expression not being well-typed due to non-linear usage of linear resources.

3.4 Lazy let bindings

In Section 2.1, we discussed how linear resources used in let-bound expressions are only consumed when the respective binders are evaluated in the let body. Moreover, resources captured by a let-bound expression cannot be used simultaneously with the let-binder in its body, since those resources would end up being consumed more than once and violate the semantics of linearity. Either the binding or the resources *must* be used.

Let-bound variables are the canonical example of a Δ -bound variable. Annotating let-bound variables with their usage environment Δ delays the consumption of resources to when the variables

themselves are used and guarantees the desired mutual exclusion property between the binder and the resources:

$$\frac{\Gamma; \Delta \vdash e : \tau \qquad \Gamma, x :_{\Delta} \tau; \Delta, \Delta' \vdash e' : \sigma}{\Gamma; \Delta, \Delta' \vdash \mathbf{let} \ x :_{\Delta} \tau = e \ \mathbf{in} \ e' : \sigma} \ (\text{Let})$$

The rule for (non-recursive) let bindings splits the linear environment in Δ and Δ' . Δ is used to type the body e of the let binding x. Perhaps surprisingly, the resources Δ are still available to type the let body e', alongside the unrestricted x binding annotated with the usage environment Δ . The resources of e being available in e' reflects how the body of a lazy binder doesn't consume resources outright, and e being an unrestricted e-bound variable ensures that using e will consume the resources e that e captures.

3.5 Recursive let bindings

Recursive let bindings also bind expressions lazily and so introduce a Δ -bound variable for each binding. The resources required to type the let-bindings are still available in the let body. These resources may be consumed via the corresponding Δ -bound variables, or directly if the let-bindings are unused.

The typing rule for recursive groups of bindings leverages both the assumption that all mutually recursive groups are strongly connected and the corollary that every binder in such a group must be typed with the same linear context, as observed in Section 2.1. Consequently, all bindings of the recursive group are introduced as Δ -bound variables sharing the same usage environment – using any one of the bindings in a recursive group entails consuming all resources required to type that same group – so, at most a single binder from the group can be used in the let body.

$$\frac{\overline{\Gamma, \overline{x_i}:_{\Delta}\tau_i}; \Delta \vdash e_i : \sigma}{\Gamma; \Delta, \Delta' \vdash \text{let rec } \overline{x_i}:_{\Delta}\tau_i = e_i \text{ in } e' : \sigma} \text{ (LetRec)}$$

This formulation is not syntax-directed since it does not describe how to annotate binders with a particular usage environment Δ . Our system simply assumes the annotations are explicitly provided in the source and checks the correctness of the overall recursive let.

In practice, determining this typing environment Δ amounts to finding a least upper bound of the resources needed to type each mutually-recursive binding that (transitively) uses all binders in the recursive group. Our implementation uses a naive $O(n^2)$ algorithm for inferring usage environments of recursive bindings. Moreover, we note that inference of usage environments for recursive binding groups bears some resemblance to the inference of principle types for recursive bindings traditionally achieved through the Hindley–Milner inference algorithm [Damas and Milner 1982]. We leave the exploration of this connection as future work.

3.6 Case Expressions

A case expression drives evaluation by evaluating its scrutinee to Weak Head Normal Form (WHNF) [Peyton Jones 1987]. Then, the case alternative whose pattern matches the result of evaluating the scrutinee is taken². An expression in WHNF can either be a λ -abstraction (λx . e) or a datatype constructor application ($K \bar{e}$). In both cases, the sub-expressions e or \bar{e} occurring in the lambda body or as constructor arguments need not be evaluated for the lambda or constructor application to be in WHNF (and they may capture linear resources).

A key detail is that when the scrutinee is already in WHNF, evaluation continues in the alternative simply by substituting the pattern variables by the scrutinee sub-expressions. No resource from the

²In our calculus, the alternatives are assumed to be exhaustive, i.e. there always exists at least one pattern which matches the scrutinee in its WHNF, so we're guaranteed to have an expression to progress evaluation.

scrutinee can be consumed since no computation happens at all. We introduce a typing judgement Γ ; $\Delta \Vdash e : \sigma > \overline{\Delta_i}$ to extract additional information from the static structure of terms in WHNF:

$$\frac{(\text{WHNF}_K)}{\Gamma; \cdot \vdash e_{\omega} : \tau_i} \frac{\Gamma; \Delta_j \vdash e_j : \tau_j}{\Gamma; \Delta \Vdash K \ \overline{e_{\omega} e_j} : \sigma > \overline{\Delta_j}} \frac{K : \overline{\tau_i \to_{\omega} \tau_j \multimap \sigma}}{\Gamma; \Delta \Vdash \lambda x. \ e : \sigma} \frac{(\text{WHNF}_{\lambda})}{\Gamma; \Delta \vdash \lambda x. \ e : \sigma > \Delta}$$
s judgement differs from the main typing judgement in that (1) it only applies to expression

This judgement differs from the main typing judgement in that (1) it only applies to expressions in weak head normal form, and (2) it "outputs" (to the right of >) a disjoint set of linear environments $(\overline{\Delta_i})$, where each environment corresponds to the linear resources used by a sub-expression of the WHNF expression. To type a constructor application K $\overline{e_\omega}e_i$, where e_ω are the unrestricted arguments and e_i the linear arguments of the constructor, we split the resources Δ into a disjoint set of resources $\overline{\Delta_i}$ required to type each linear argument individually and return exactly that split of the resources; the unrestricted e_ω expressions must be typed on an empty linear environment. A lambda expression is typed with the main typing judgement and trivially "outputs" the whole Δ environment, as there is always only a single sub-expression in lambda abstractions.

3.6.1 Discriminating by WHNF-ness. The dichotomy between evaluation (and resource usage) of a case expression whose scrutinee is in weak head normal form and one whose scrutinee is not leads to one of the key insights of λ_{Δ}^{π} : accurate typing of case expressions depends on whether the scrutinee is in WHNF. When the scrutinee is already in WHNF, the resources are unused upon evaluation and thus available in the alternatives. When it is not, resources will be consumed and cannot be used in the alternative. To illustrate, consider the following case expressions:

(1)
$$\lambda x$$
. case K x of $z \{_ \rightarrow x\}$ (2) λx . case f r e x of $z \{_ \rightarrow x\}$

The first function uses *x* linearly, but the second does not. Alternatives may also use the case binder or pattern variables, referring to, respectively, the whole scrutinee and all its used resources; or constructor arguments and the resources used to type them. The reader may wonder about the significance of case expressions where the scrutinee is in WHNF. We note that in an optimizing compiler (such as GHC) many intermediate programs that arise due to program transformations commonly scrutinize expressions in WHNF.

There are three competing ways to use the resources from a scrutinee in WHNF in a case alternative: directly, via the case binder, or by using pattern-bound variables. Recall how Δ -bound variables encode mutual exclusion between alternative ways of consuming resources – it follows that case binders and pattern-bound variables are yet another instance of Δ -bound variables. This suggests the following first rule for cases of scrutinees already in WHNF:

$$\frac{(\mathsf{CASE}_{\mathsf{WHNF}})}{\mathsf{e} \; \mathsf{is} \; \mathsf{in} \; \mathsf{WHNF}} \qquad \Gamma; \Delta \Vdash e : \tau > \overline{\Delta_i} \qquad \overline{\Gamma, z :_{\overline{\Delta_i}} \tau; \overline{\Delta_i}, \Delta' \vdash \rho \; \Rightarrow \; e' :_{\overline{\Delta_i}}^z \sigma}}{\Gamma; \Delta, \Delta' \vdash \mathsf{case} \; e \; \mathsf{of} \; z :_{\overline{\Delta_i}} \tau \; \{ \overline{\rho} \; \Rightarrow \; e' \} : \sigma}$$

First, we assert this rule is only applicable to expressions in weak head normal form. Second, we appeal to the typing judgement for expressions in WHNF to determine the split of resources amongst the scrutinee sub-expressions. Finally, we type all case alternatives with the same context, using the introducing the case binder z in the unrestricted environment as a Δ -bound variable whose usage environment is the linear resources $\overline{\Delta_i}$ used to type the scrutinee. Those same resources $\overline{\Delta_i}$ are again made available in the linear typing environment of the alternative, similarly to how the resources used to type a *Let* binder are still available in the continuation. Although the main idea for typing "WHNF case expressions" is conveyed by this rule, the full rule of our system is slightly more involved, as will be seen after discussing cases of scrutinees not in WHNF.

The alternative judgement Γ ; $\Delta \vdash \rho \Rightarrow e :_{\Delta}^{z} \sigma$ is used to type case alternatives, encompassing three modes that are distinguished by the kind of arrow that is used: for alternatives of case expressions whose scrutinee is in WHNF ($\Rightarrow_{\text{WHNF}}$), not in WHNF ($\Rightarrow_{\text{NWHNF}}$), and for alternatives agnostic to the WHNF-ness of the scrutinee (\Rightarrow), with \Rightarrow also generalizing the other two.

First, an alternative whose pattern is a constructor with n > 0 linear components under the $\Rightarrow_{\text{WHNF}}$ mode is typed as:

$$\frac{(\text{AltN}_{\text{WHNF}})}{\Gamma, \overline{x}:_{\omega} \overline{\tau}, \overline{y}_{i}:_{\Delta_{i}} \overline{\tau_{i}}^{n}; \Delta \vdash e : \sigma}}{\Gamma; \Delta \vdash K \ \overline{x}:_{\omega} \overline{\tau}, \overline{y}_{i}:_{1} \overline{\tau_{i}}^{n}} \Rightarrow_{\text{WHNF}} e :_{\overline{\Delta_{i}}^{n}} \sigma}$$

The rule states that, for such a match on a scrutinee already in WHNF, we introduce the linear components of the pattern as Δ -bound variables with usage environment matching the linear resources required to type the corresponding constructor argument in the scrutinee. The resources required to type each constructor sub-expression come annotated in the judgement (as : $_{\overline{\Delta_i}}$). Unrestricted fields of the constructor are introduced as unrestricted variables. We note that the typing environment Δ always contains the resources $\overline{\Delta_i}$ in uses of the alternative judgement.

Second, the rule for alternatives that match on the wildcard pattern, regardless of the WHNF status of the scrutinee (note the use of \Rightarrow , which is applicable both under \Rightarrow_{WHNF} and \Rightarrow_{NWHNF}):

$$\frac{\Gamma; \Delta \vdash e : \sigma}{\Gamma; \Delta \vdash _ \Rightarrow e :_{\Delta_{s}}^{z} \sigma}$$

To type a wildcard alternative we simply type the expression with the main judgement, ignoring all annotations. The case binder will have already been introduced in the environment with the appropriate usage environment by the relevant case expression rule.

Lastly, consider an alternative matching on a data constructor without any linear components. The linear resources used to type a scrutinee matching such a pattern are fully consumed during evaluation: the resulting unrestricted constructor application can only be well-typed under an empty linear environment. Consequently, the resources to type the scrutinee which are carried over to the case-alternative environments by the $Case_{WHNF}$ and $Case_{NWHNF}$ rules (see the latter below), must be reactively destroyed from branches where the pattern is unrestricted. In failing to do this, we would allow the fully-consumed resources to be used again in the branch. The Alt0 rule essentially encodes this insight, and is applicable regardless of the WHNF status of the scrutinee (\Rightarrow mode), as long as the constructor pattern has no linear fields:

$$\frac{\Gamma\left[\cdot/\Delta_{s}\right]_{z}, \overline{x}:_{\omega}\tau; \Delta\left[\cdot/\Delta_{s}\right] + e : \sigma}{\Gamma; \Delta \vdash K \ \overline{x}:_{\omega}\tau \ \Rightarrow \ e :_{\Delta_{s}}^{z} \sigma}$$

The rule deletes the annotated scrutinee environment Δ_s from the linear typing context (written $\Delta[\cdot/\Delta_s]$, a substitution of the scrutinee typing environment by the empty linear environment ·) and from the usage environment of the case binder z, written $\Gamma[\cdot/\Delta_s]_z$ to delete Δ_s from the usage environment of the Δ -bound variable z in Γ . Given that the case-binder z is always annotated with Δ_s (essentially, the linear environment typing the scrutinee), this latter substitution makes the case-binder unrestricted in branches with unrestricted patterns.

This rule observes that any linear resources required to type an expression that reduces to an unrestricted constructor application are fully consumed through evaluation and the resulting unrestricted expression can be freely discarded or duplicated (e.g. via the case binder).

- 3.6.2 Irrelevant resources. When the scrutinee is not in WHNF we cannot statically determine which resources are consumed upon scrutinee evaluation or which are transferred into the branch continuation e.g. as in case $f \times f = \{x \in A \in A\}$:
 - (1) Cases with non-WHNF scrutinees must be typed such that no resource from the scrutinee can be used directly in the branch again, or we risk consuming the same resource twice. In the example, *x* must not directly occur in *e* since *f* could have consumed it.
 - (2) On the other hand, since evaluation to WHNF does not guarantee all resources are consumed (unless the result matches an unrestricted pattern), we must ensure that any resources that transfer from the scrutinee to the alternative are ultimately consumed (in the example, x is not consumed by evaluating the scrutinee for f = K).

Considering (1), we must resort to a more uniform method of guaranteeing we "finish consuming" the scrutinee. There are only two ways of uniformly referring to the remaining linear resources in the newly-evaluated scrutinee:

- The case binder, which refers to the whole result of evaluating the scrutinee;
- The linear components of a constructor pattern, which refer to expressions that may contain linear resources.

To encode linear resources that cannot be directly used (i.e., to which the VAR rule is not applicable) we introduce *irrelevant* resources, written as linear resources within square brackets $[x:\sigma]$ and pointwise lifted to contexts and usage environments as $[\Delta]$. Irrelevant resources are linear resources in every sense, meaning they must be used exactly once. However, since irrelevant resources cannot be discarded or used directly, they have to be consumed *indirectly* via Δ -bound variables, namely, the case binder, or, in mutual exclusion, the linear pattern-bound variables.

Hence, to type a case expression whose scrutinee is not in WHNF, we type the scrutinee with linear resources Δ and the case alternatives by introducing the case binder with an irrelevant usage environment $[\Delta]$, having the same irrelevant $[\Delta]$ in the typing environment, and annotating the judgement with the irrelevant resources for use in the \Rightarrow_{NWHNF} judgement:

$$\frac{\text{(Case}_{\text{Not WHNF}})}{\text{e is not in } WHNF} \qquad \Gamma; \Delta \vdash e : \tau \qquad \overline{\Gamma, z_{:[\Delta]}\tau; [\Delta], \Delta' \vdash \rho \Rightarrow_{\text{NWHNF}} e' :_{[\Delta]}^z \sigma}{\Gamma; \Delta, \Delta' \vdash \mathbf{case} \ e \ \mathbf{of} \ z_{:[\Delta]}\tau \ \{ \overline{\rho \Rightarrow e'} \} : \sigma}$$

Note how the rule is similar to the one for scrutinees in WHNF, but requires the resources in the case binder, typing environment, and judgement annotation to be made "irrelevant".

Finally, we recall the tentative $Case_{WHNF}$ rule presented before and highlight its flaw: the Γ ; Δ $\vdash \rho \Rightarrow_{WHNF} e :_{\Delta s}^z \sigma$ judgement is only well-defined for patterns ρ matching the WHNF form of the scrutinee, as the distribution of resources per constructor components only makes sense for the constructor pattern matching the scrutinee. Alternatives not matching the scrutinee could use resources arbitrarily as they will never be executed. We uniformly treat non-matching alternatives as if the scrutinee were not in WHNF. Having introduced irrelevant resources, we can now present the full $Case_{WHNF}$ rule:

$$\begin{array}{c} \text{(CASE}_{\text{WHNF}}) \\ \text{e is in } \textit{WHNF} \quad e \text{ matches } \rho_j \quad \overline{\Gamma, z_{: [\Delta]} \tau; [\Delta], \Delta' \vdash \rho \Rightarrow_{\text{NWHNF}} e' :_{[\Delta]}^z \sigma } \\ \hline \qquad \Gamma; \Delta \Vdash e : \tau \rhd \overline{\Delta_i} \quad \Gamma, z_{: \overline{\Delta_i} \tau}; \overline{\Delta_i}, \Delta' \vdash \rho_j \Rightarrow e' :_{\overline{\Delta_i}}^z \sigma \\ \hline \qquad \Gamma; \Delta, \Delta' \vdash \mathbf{case} \ e \ \mathbf{of} \ z_{: \overline{\Delta_i} \tau} \tau \ \{ \overline{\rho \Rightarrow e'} \} : \sigma \end{array}$$

where we appeal to the pattern judgment for WHNF scrutinees for pattern ρ_j , allowing the $\overline{\Delta_i}$ resources to be consumed either through the case-binder or outright; and to the non-WHNF

judgment for the remaining branches, where the Δ resources may only be consumed via the case-binder. We again note that it might seem unusual to specialize a rule for expressions in WHNF, as programs scrutinizing an expression in WHNF are rarely written by humans. Yet, our system is designed to be suitable for optimising compilers in which intermediate programs commonly scrutinize expressions in WHNF.

3.6.3 Splitting and tagging fragments. As opposed to scrutinees in WHNF, where resources can be cleanly divided amongst the various sub-expressions of a constructor application and, henceforth, each pattern variable, there is no direct mapping between the resources typing a scrutinee not in WHNF and the usage environments of pattern variables in any alternative.

We introduce *tagged resources* as a means to enforce that all pattern-bound variables for a scrutinee not in WHNF are either *jointly* used to consume all resources occurring in the environment, or not at all (instead, the case binder may be used). Given irrelevant resources $[\Delta_s]$ used to type a scrutinee, and a pattern $K[\overline{x_\omega}, \overline{y_i}]$ with i linear components, we assign a usage environment Δ_i to each linear pattern variable where Δ_i is obtained from the scrutinee environment *tagged* with the constructor name and linear-variable index $[\Delta_s] \# K_i$. Then, $y_i : \Delta_i \sigma$ is introduced in Γ , just like other Δ -bound variables.

$$\frac{(\text{AltN}_{\text{Not WHNF}})}{\Gamma, \overline{x}:_{\omega} \overline{\tau}, \overline{y}_{i}:_{\Delta_{i}} \overline{\tau}_{i}; \Delta \vdash e : \sigma} \qquad \overline{\Delta_{i}} = \overline{\Delta_{s} \# K_{i}}^{n} \qquad \frac{(\text{Split})}{\Gamma; \Delta, x :_{1} \tau \vdash e : \tau} \qquad K \text{ has } n \text{ linear arguments}}{\Gamma; \Delta, \overline{x}:_{1} \overline{\tau} \# K_{i}}^{n} \mapsto_{\text{NWHNF}} e :_{\Delta_{s}}^{z} \sigma}$$

Having uniquely tagged the resources in the usage environment of each pattern variable, we need only express that (1) the pattern Δ -bound variables can be used (i.e., the tagged resources need be available in the linear environment in order for the Var $_{\Delta}$ rule to be applicable) and (2) if a K-pattern Δ -bound variable is consumed, all remaining K- Δ -bound variables, standing for the remaining linear components of the same pattern, must also be consumed.

To satisfy these two constraints, we introduce a Split rule that allows a linear resource $x_{1}\sigma$ to be split into n resources at a given constructor K, where n is the number of linear components of the constructor, and each resource resulting from the split is tagged with K and the positional index of a linear component (i.e. $x_{1}\sigma$ can be split into $\overline{x_{1}\sigma\#K_{i}}^{n}$). By assigning to each pattern variable a fragment of the scrutinee resources (with a tag), we require the scrutinee resources to be Split in order to use pattern variables at all. Moreover, the remaining tagged resources cannot be used directly, yet need to still be consumed exactly once. Thus, the choice to consume the tagged resources via the usage environments of the other pattern variables is forced, ensuring that no variable for a constructor's linear component can go unused (the tagged environments are disjoint). Tagged resources are inspired by fractional permissions in separation logic [Boyland 2003].

For instance, in the term λx . case f x of z {K a $b \to (a,b)$ }, where x is a linear variable, the case alternative is typed with [x], the case binder z is introduced as $z:_{[x]}\tau$, and the pattern variables are introduced as $a:_{[x]\#K_1}\tau$ and $b:_{[x]\#K_2}\tau$, assuming both components of K are linear. The occurrences of a and b are well-typed because we can first Split $[x:_1\sigma]$ into $[x:_1\sigma]\#K_1$, $[x:_1\sigma]\#K_2$, noting that Split can be applied both to relevant and proof irrelevant linear resources in Δ .

4 Metatheory of Linear Core

We develop the metatheory of Linear Core by first presenting the operational semantics of λ_{Δ}^{π} , consisting of a lazy natural semantics in the style of Launchbury [Launchbury 1993]. This semantics conveys the expected behavior of an implementation of lazy evaluation and is agnostic to any linearity information. Such a semantics is ill-suited for reasoning about linearity and so we develop an *instrumented* linearity-aware semantics that is enriched with sufficient information to allow us

$$\frac{\Theta: e \Downarrow \Theta': \Lambda p.e' \quad \Theta': e'[\pi/p] \Downarrow \Theta'': v}{\Theta: e \pi \Downarrow \Theta'': v}$$

$$\frac{\Theta: e \Downarrow \Theta': \lambda y :_{\pi} \sigma.e' \quad \Theta': e'[x/y] \Downarrow \Theta'': v}{\Theta: e \pi \Downarrow \Theta'': v}$$

$$\frac{\Theta: e \Downarrow \Theta': \lambda y :_{\pi} \sigma.e' \quad \Theta': e'[x/y] \Downarrow \Theta'': v}{\Theta: e x \Downarrow \Theta'': v}$$

$$\frac{(\Theta, x :_{\omega} \sigma = e) : e \Downarrow \Theta': v}{(\Theta, x :_{\omega} \sigma = e) : x \Downarrow (\Theta', x :_{\omega} \sigma = v) : v} \qquad \frac{(\Theta, x :_{\omega} \sigma = e) : e' \Downarrow \Theta': v}{\Theta: \text{let } x :_{\Delta} \sigma = e \text{ in } e' \Downarrow \Theta': v}$$

$$\frac{\Theta: e \Downarrow \Theta': K \overline{x_i} \parallel \Theta: K \overline{x_i}}{\Theta: \text{case } e \text{ of } z :_{\Delta'} \sigma \{ \dots, K \overline{y_i} \Rightarrow e', \dots \} \Downarrow \Theta'': v}$$

$$\frac{\Theta: e \Downarrow \Theta': K \overline{x_i} \mid \Theta': e'[K \overline{x_i}/z] \Downarrow \Theta'': v}{\Theta: \text{case } e \text{ of } z :_{\Delta'} \sigma \{ \dots, K \overline{y_i} \Rightarrow e', \dots \} \Downarrow \Theta'': v}$$

$$\frac{\Theta: e \Downarrow \Theta': K \overline{x_i} \mid \Theta': e'[K \overline{x_i}/z] \Downarrow \Theta'': v}{\Theta: \text{case } e \text{ of } z :_{\Delta'} \sigma \{ \dots, K \overline{y_i} \Rightarrow e', \dots \} \Downarrow \Theta': v}$$

Fig. 6. Natural Semantics of λ_{Λ}^{π}

to establish type safety, thus showing that linearity is preserved. We show that the natural and instrumented semantics are bisimilar for well-typed terms and so derive type safety in the natural semantics (Section 4.2). Finally, we show that multiple optimising transformations such as those performed by GHC on Linear Core are type preserving (Section 4.3), and thus linearity preserving.

4.1 Operational Semantics

The Launchbury-style natural semantics, presented in Figure 6, captures standard lazy evaluation while ignoring linearity information. The natural semantics is equipped with an evaluation environment that maps variables to expressions and is mutated in order to express shared evaluation. Terms are explicit about their sharing (in the sense of shared reductions in lazy evaluation), and so we perform a translation step that makes all sharing explicit through let-binders (see Appendix). We write let $\overline{x_i}$: $\overline{x_i}$ or $\overline{e_i}$ in e' for the iterated let-binding of variables $\overline{x_i}$.

The natural semantics is defined via the relation $\Theta: e \Downarrow \Theta': v$ where e is an expression, v a value, and Θ and Θ' are evaluation environments with bindings of the form $x:_{\omega} \sigma = e$ which assigns the expression e to the variable e of the given type e. The rules are standard, augmenting the runtime environment with let-bound expressions which are evaluated (only once) if the corresponding let-bound variable is forced.

While the natural semantics is a mostly direct model of an implementation of lazy evaluation, it is inconvenient for formal reasoning about well-typed linear terms. We follow an approach similar to that of [Bernardy et al. 2017a,b], defining an *instrumented* operational semantics in which linear variables are erased from the runtime environment once they are forced, ensuring that any term violating linearity will result in a stuck state. The semantics is also type-aware and carries with it sufficient data to reconstruct typing derivations for the purposes of showing type safety.

The instrumented semantics is presented in Figure 7, defining the judgment Γ ; $\Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid v) : \sigma, \Sigma$ where Γ and Δ are typing contexts for expressions in Σ (Γ tracks unrestricted and Δ -bound variables, Δ tracks linear assumptions); Θ and Θ' are evaluation environments, consisting of bindings of the form $x :_{\pi} \sigma = e$ or $x :_{\Delta} \sigma = e$; e is the expression being evaluated and v its resulting value, both of type σ ; Σ is a list of assignments of the form $e :_{\tau}$ which are expressions in which bindings in Θ may also be used. This additional data is needed to inductively show that the overall evaluation state is well-typed (and linearity preserving). The rules mostly mirror those of the natural semantics, with the exception of the linear variable rule which erases the binding from

$$\begin{array}{c} \Gamma;\Delta \vdash (\Theta \mid \Delta p.e) \Downarrow (\Theta \mid \Delta p.e) : \forall p.\sigma,\Sigma \\ \Gamma;\Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid \Delta p.e') : \forall p.\sigma,\Sigma \quad \Gamma;\Delta \vdash (\Theta' \mid e'[\pi/p]) \Downarrow (\Theta'' \mid v) : \sigma[\pi/p],\Sigma \\ \hline \Gamma;\Delta \vdash (\Theta \mid e \pi) \Downarrow (\Theta'' \mid v) : \sigma[\pi/p],\Sigma \\ \hline \Gamma;\Delta \vdash (\Theta \mid kx :_{\pi} \sigma.e) \Downarrow (\Theta \mid kx :_{\pi} \sigma.e) : \sigma \rightarrow_{\pi} \tau,\Sigma \\ \hline \Gamma;\Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid ky :_{1} \sigma.e') : \sigma \rightarrow_{1} \tau,x : \sigma,\Sigma \quad \Gamma;\Delta \vdash (\Theta',y :_{1} \sigma = x \mid e') \Downarrow (\Theta'' \mid v) : \tau,\Sigma \\ \hline \Gamma;\Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid ky :_{\omega} \sigma.e') : \sigma \rightarrow_{\omega} \tau,x : \sigma,\Sigma \quad \Gamma;\Delta \vdash (\Theta' \mid e'[x/y]) \Downarrow (\Theta'' \mid v) : \tau,\Sigma \\ \hline \Gamma;\Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid ky :_{\omega} \sigma.e') : \sigma \rightarrow_{\omega} \tau,x : \sigma,\Sigma \quad \Gamma;\Delta \vdash (\Theta' \mid e'[x/y]) \Downarrow (\Theta'' \mid v) : \tau,\Sigma \\ \hline \Gamma;\Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid v) : \sigma,\Sigma \quad \Delta'' = \Delta' \setminus (\Theta \setminus \Theta') \\ \hline \Gamma;\Delta \vdash (\Theta,x :_{\Delta'} \sigma = e \mid e) \Downarrow (\Theta' \mid v) : \sigma,\Sigma \quad \Delta'' = \Delta' \setminus (\Theta \setminus \Theta') \\ \hline \Gamma;\Delta \vdash (\Theta,x :_{\Delta'} \sigma = e \mid x) \Downarrow (\Theta',x :_{\Delta''} \sigma = v \mid v) : \sigma,\Sigma \\ \hline \Gamma;\Delta \vdash (\Theta,x :_{\Delta'} \sigma = e \mid x) \Downarrow (\Theta' \mid v) : \tau,\Sigma \\ \hline \Gamma;\Delta \vdash (\Theta,x :_{\Delta'} \sigma = e \mid e') \Downarrow (\Theta' \mid v) : \tau,\Sigma \\ \hline \Gamma;\Delta \vdash (\Theta \mid kx :_{\Delta'} \sigma = e \mid e') \Downarrow (\Theta' \mid v) : \tau,\Sigma \\ \hline \Gamma;\Delta \vdash (\Theta \mid kx :_{\Delta'} \sigma = e \mid e') \Downarrow (\Theta' \mid v) : \tau,\Sigma \\ \hline \Gamma;\Delta \vdash (\Theta \mid kx :_{\Delta'} \sigma = e \mid e') \Downarrow (\Theta' \mid v) : \tau,\Sigma \\ \hline \Gamma;\Delta \vdash (\Theta \mid kx :_{\Delta'} \sigma = e \mid e') \Downarrow (\Theta' \mid v) : \tau,\Sigma \\ \hline \Gamma;\Delta \vdash (\Theta \mid exee e of z :_{\Delta'} T \{\dots,K y_{i,j} \Rightarrow e',\dots\}\} \Downarrow (\Theta'' \mid v) : \tau,\Sigma \\ \hline \Gamma;\Delta \vdash (\Theta \mid exee e of z :_{\Delta'} T \{\dots,K y_{i,j} \Rightarrow e',\dots\}\} \Downarrow (\Theta'' \mid v) : \tau} \\ \hline \Gamma;\Delta \vdash (\Theta \mid exee e of z :_{\Delta'} T \{\dots,L w_{i,j} \Rightarrow e',\dots\}\} \Downarrow (\Theta'' \mid v) : \tau \\ \hline \Gamma;\Delta \vdash (\Theta \mid exee e of z :_{\Delta'} T \{\dots,L w_{i,j} \Rightarrow e',\dots\}\} \Downarrow (\Theta'' \mid v) : \tau$$

Fig. 7. Instrumented Operational Semantics

the execution environment. Both semantics define evaluation relations of the form $a \downarrow b$, where a and b are evaluation states. We also rely on a notion of partial derivations and partial evaluation (see Appendix), written $a \downarrow b$, which allows us to state and reason about a progress property.

4.2 Type Safety

Our results rely on a series of lemmas relating linear, unrestricted and Δ -bound variables, as well as a lemma characterizing irrelevant resources.

 Δ -bound variables. A well-typed program with a linear variable $(x:_1\tau)$ is equivalently well-typed if the linear variable were instead Δ -bound $(x:_{\Delta}\sigma)$ with usage environment Δ , with Δ available in the linear context and with any occurrences of x in usage environments in Γ expanded into x's usage environment Δ , with the abuse of notation $\Gamma[\Delta/x]$.

Lemma 4.1 (Linear to
$$\Delta$$
-bound). If Γ ; Δ' , x : $_1\tau \vdash e : \sigma$ then $\Gamma[\Delta/x]$, x : $_{\Delta}\sigma$; Δ , $\Delta' \vdash e : \sigma$ (Δ fresh).

Dually, any Δ -bound variable with an irrelevant usage environment can be made into a linear variable as long as its (irrelevant) environment is dropped from the linear context and all the occurrences of the set of variables $[\Delta]$ in usage environments in Γ are collapsed to x, written $\Gamma[x/[\Delta]]$.

Lemma 4.2 (Δ-bound to Linear). If
$$\Gamma$$
, $x:_{[\Delta]}\tau$; $[\Delta]$, $\Delta' \vdash e : \sigma$ then $\Gamma[x/[\Delta]]$; Δ' , $x:_1\tau \vdash e : \sigma$.

Finally, unrestricted and Δ -bound variables with an empty usage environment are equivalent.

Lemma 4.3 (Unrestricted and Δ-bound).
$$\Gamma, x:_{\omega} \tau; \Delta \vdash e : \sigma \ iff \Gamma, x:_{\emptyset} \tau; \Delta \vdash e : \sigma$$

Irrelevance. Irrelevant resources (Section 3.6.2) can only be consumed indirectly, essentially encoding that the non-WHNF case scrutinee resources must be consumed through the case binder or the linear pattern-bound variables. As a case expression is evaluated, the scrutinee will eventually reduce to WHNF, which must then be typed with rule $Case_{WHNF}$. Crucially, the two case rules must be in harmony in the system, in the sense that case expressions typed using the $Case_{Not\ WHNF}$ rule must also be well-typed by the $Case_{WHNF}$ once the scrutinee is evaluated to WHNF.

Lemma 4.4 (Irrelevance). If
$$\Gamma$$
, $z:_{[\Delta]}\tau$; $[\Delta]$, $\Delta' \vdash \rho \Rightarrow_{\mathsf{NWHNF}} e:_{[\Delta]}^z \sigma$ then Γ , $z:_{\Delta^\dagger}\tau$; Δ^\dagger , $\Delta' \vdash \rho \Rightarrow e:_{\Delta^\dagger}^z \sigma$, for any Δ^\dagger

The *Irrelevance* lemma shows that if a case alternative is well-typed with irrelevant resources, then it is well-typed with arbitrary resources. However, typing a case alternative with irrelevant resources is not complete wrt using arbitrary resources – a counter example needs only to use a resource directly.

We prove type safety of the Linear Core system via the standard type preservation and progress results. As is customary, we make use of multiple substitution lemmas, one for each kind of variable: unrestricted variables $x:_{\Delta}\tau$, linear variables $x:_{1}\tau$, and Δ -bound variables $x:_{\Delta}\sigma$. The development is reported in the Appendix, where type safety is established directly for the instrumented semantics and ported to the natural semantics via a bisimulation argument.

We adopt the standard semantic technique of dealing with badly-formed terms, such as trying to add two values that are not numbers, by having such terms be stuck. In our setting, linear variables used more than once result in stuck states. Thus, progress ensures that linear functions consume their arguments at most once if their result is consumed exactly once. Type preservation ensures that evaluation of a closed term returns an evaluation environment with no linear variables.

Theorem 4.5 (Type Preservation). For any well-typed evaluation state a, if $a \parallel b$ or $a \parallel^* b$ then b is well-typed.

Theorem 4.6 (Progress). Let a be a well-typed evaluation state. For any partial evaluation a $\downarrow ^* b$ the evaluation can be extended.

4.3 Optimisations preserve linearity

A main goal of the Linear Core type system is to serve as a typed intermediate representation for optimising compilers of lazy languages with linearity. In light of this, we show that multiple optimising transformations are type preserving in Linear Core.

We describe Core-to-Core transformations as $e_1 \implies e_2$, where e_1 is an arbitrary well-typed expression matching a certain shape (or of a certain type), which is transformed into expression e_2 . Validating a transformation entails showing that e_2 is well-typed. We often annotate the arrow with the name of the transformation when describing chains of such transformations.

Fig. 8. Optimising transformations validated by Linear Core (for commuting case-let we require e_1 not to capture the case binder, variables bound by the pattern ρ and the context E[-] – see Appendix for details).

The list of transformations we have validated is given in Figure 8. The behaviour of most transformations can be inferred from its name: Inlining substitutes a let-bound expression in the let body; β -reduction and case of known constructor effectively perform evaluation of the given expression; case-of-case commutes cases in scrutinee position inside the outer case alternatives; the five commuting conversions allow let-bound expressions to be commuted with other Core constructs; η -expansion and reduction apply the η -laws to terms of function type; finally, binder swap substitutes occurrences of a scrutinee variable for the case binder. While many of the transformations above are not obviously optimising, they are chained in the compiler pipeline to expose optimisation opportunities.

Reverse Binder Swap. Reverse binder swap is dual to the binder swap transformation, substituting occurrences of the case binder z by the scrutinee in the special case when the scrutinee is a variable x. By using the scrutinee x instead of the case binder, we might be able to float out expressions from the alternative using the case binder:

case x of
$$z \{ \rho \to e \} \Rightarrow \text{case } x \text{ of } z \{ \rho \to e[x/z] \}$$

Even though GHC applies the reverse binder swap transformation in the Core-to-Core passes, this optimisation violates linearity in Linear Core because the transformed program is rejected – the scrutinee x, a single variable, is not considered to be in WHNF by our system, thus occurrences of x in any alternative are ill-typed in it. This is a conservative choice in λ_{Δ}^{π} for typing the subtle case where the scrutinee is a single variable.

Consider **case** x **of** $_\to x$. Perhaps surprisingly, this expression can be considered semantically linear under call-by-need but violates linearity under call-by-name: Under call-by-need, if x refers to an unevaluated expression, scrutinizing it forces the expression bound by x to WHNF. In a subsequent use of x in an alternative, x refers to the already evaluated scrutinee in WHNF. Since x

is just another name (like the case binder) for the scrutinee in WHNF, its use is also valid in the alternatives. Dually, if x refers to an expression already in WHNF, no evaluation takes place and x is still just another name for the scrutinee in WHNF.

However, under call-by-name, evaluation of *x* is not shared:

$$(\lambda x. \mathbf{case} \ x \mathbf{of} \ _ \to x) \ (use \ y) \Longrightarrow_{\mathsf{CBName}} \mathbf{case} \ use \ y \mathbf{of} \ _ \to use \ y$$

After the reduction, y is a linear variable consumed both in the scrutinee and in the case alternative – y has been duplicated! If the application is instead reduced call-by-need, y is not duplicated:

$$(\lambda x. \mathbf{case} \ x \mathbf{of} \ _ \to x) \ (use \ y) \Longrightarrow_{\mathsf{CBNeed}} \mathbf{let} \ x = use \ y \mathbf{in} \mathbf{case} \ x \mathbf{of} \ _ \to x$$

In an optimising compiler such as GHC, there is no definitive evaluation strategy. It is mostly up to heuristics to determine what to inline \grave{a} la call-by-name and what to create let-bindings for \grave{a} la call-by-need. In this flexible evaluation setting, we highlight the reverse binder swap is only a linearity-preserving transformation as long as a linear variable with more than one occurrence is never substituted for an expression with free linear variables.

5 Linear Core as a GHC Plugin

We implemented Linear Core as a plugin for the Glasgow Haskell Compiler. GHC plugins allow developers to inspect and modify programs being compiled by GHC, at different stages of compilation [Pickering et al. 2019]. The Linear Core GHC plugin [Mesquita and Toninho 2025] is a typechecker. For every single intermediate Core program produced by GHC, our plugin checks the linearity of all expressions, failing if a linear resource is not used *exactly once* according to our typing discipline. The plugin implementation of Linear Core provides us with strong confidence that the Linear Core type system is suitable for the intermediate language of an optimizing compiler, since it is successful in accepting the vast majority of real-world linear programs produced by the GHC optimizer, and even identified some which violate linearity.

The implementation of a Linear Core checker does not follow directly from the type system because Linear Core is not syntax-directed (the design aims to be as simple to reason with as possible). Specifically, the features of Linear Core that are not syntax-directed are essentially the splitting of linear resources amongst sub-derivations and consuming fragments of resources through pattern-bound variables. As standard in substructural typing systems, our rules non-deterministically split linear resources as needed. In our implementation, we thread input/output linear resources through each derivation [Cervesato et al. 2000]. To use pattern variables with usage environment $[\Delta] \# K_i$, the scrutinee resources must first be Split into fragments. Instead of guessing which resources need to be split, the implementation consumes tagged fragments of a resource as needed, i.e., only when consuming a resource tagged K_j do we Split that resource on K and consume the K_j fragment. Finally, our implementation must infer the usage environments of binders in a recursive let group before using them to typecheck the let body. We use a naive $O(n^2)$ algorithm (where n is the number of let bindings) to determine these usage environments. We leave the study of the soundness and completeness of this implementation strategy to future work.

The results of running the Linear Core GHC plugin on real-world libraries with linear types are reported in Figure 9. We compiled the Haskell standard library for programming with linear types, linear-base (4000 lines, comprising over 100 modules), linear-smc [Bernardy and Spiwack 2021] (1500 lines), priority-sesh [Kokke and Dardha 2021] (1400 lines), text-builder-linear [Lelechenko 2024] (1800 lines), and linear-generics [Spiwack et al. 2025] (2000 lines), using our plugin. We count the number of programs accepted by our implementation, where each top-level binding in a module counts as a program, and every such binding is checked once per optimisation pass, i.e. we check all intermediate programs produced by GHC. The total of rejected programs is

counted in the same way, without halting compilation. We note that our implementation is not performance conscious and types every intermediate program from scratch.

Library	Total Accepted	Total Rejected	Unique Rejected	Accept Rate
linear-base	48018	735	51	98.50%
linear-smc	18759	5	3	99.97%
priority-sesh	760	95	6	88.89%
text-builder-linear	4858	34	6	99.30%
linear-generics	63332	0	0	100.00%

Fig. 9. Linear Core Plugin on Linear Libraries with GHC 9.10.3

The results indicate that our mostly direct implementation of Linear Core is successful in accepting the vast majority of the thousands of intermediate programs produced by GHC when compiling libraries that make extensive use of linear types. However, Linear Core does not accept every single program produced by the optimisation pipeline of GHC. Our manual analysis of rejected programs revealed excellent results. Programs considered ill-typed by the Linear Core implementation include programs that:

- we expected GHC to produce, but know are not seen as linear by λ_{Δ}^{π} . Namely, programs resulting from the reverse binder swap which scrutinize a variable and then use it again in the case alternatives. As discussed in Section 4.3, these programs could be understood as linear as long as all applications of linear functions are reduced call-by-need
- were rejected due to the lack of "linearity coercions" in λ_{Δ}^{π} . Extending Linear Core with these coercions would make these programs be accepted (as discussed in Section 6.1);
- actually violate linearity. These are the programs we ultimately want to spot and reject by
 having a linearly typed Core validate intermediate programs. We found both a program which
 outright discarded a linear resource and that rewrite rules implementing stream fusion [Coutts
 et al. 2007; Gill et al. 1993] (among other rules) could produce linearly-invalid programs (in
 one such case, the type of build was not general enough to accommodate linearity).

Besides validating that our implementation is faithful to the λ_{Δ}^{π} system, these programs provide insight as to what more is needed to fully understand linearity in a mature optimising compiler.

6 Related and Future Work

Linear Haskell. Linear Haskell [Bernardy et al. 2017b] augments the Haskell surface language with linear types, but it is not concerned with extending GHC's intermediate language(s), which presents its own challenges. However, the implementation of Linear Haskell in GHC does modify and extend Core with linearity/multiplicity annotations. Core's type system is unable to type semantic linearity of programs in the sense elaborated in our work, which results in Core rejecting most linear programs resulting from optimising transformations that leverage the non-strict semantics of Core. Our work overcomes the limitations of Core's linear type system derived from Linear Haskell by understanding linearity semantically in the presence of laziness, and by showing that multiple Core-to-Core optimisations employed by GHC are linearity preserving, using a linearity-aware semantics that is essentially identical to that of [Bernardy et al. 2017a]. Linear Core can also be seen as a system that validates the programs written in Linear Haskell and are compiled by GHC, by guaranteeing (through typing) that linear resources are still used exactly once throughout the optimising transformations. We note that both Linear Haskell and our work has no special treatment of exceptions. This is also the case in GHC Core, where exceptions have no special status

and so there is no natural way of dealing with the interaction between exceptions and linearity in Core itself without a major overhaul of the exception mechanisms of the language.

Linearity-aware Semantics. Several other work explore linearity aware semantics, dating back to [Chirimar et al. 1996] and [Turner and Wadler 1999]. The former uses a computational interpretation of linear logic to give an account of reference counting. The latter present a heap-based operational interpretation of linear logic and explore memory usage properties under call-by-name and call-by-need, showing both satisfy different properties under this lens. Their semantics is similar to ours, but does not account for pattern matching. More recently, [Choudhury et al. 2021; Marshall et al. 2022] and [Bernardy et al. 2017a,b] all make use of usage-aware semantics. Choudhury et al. [2021] do so in the context of graded dependent type theory, using the usage-aware heap-based operational semantics to show correct accounting of resource usage in their setting, as we do in ours, but they do not model call-by-need, only call-by-name. The work of [Marshall et al. 2022] is applies the ideas of [Choudhury et al. 2021] to a non-graded call-by-name setting. Bernardy et al. [2017b] provide a similar semantics to ours, but present only a naive linear typing system that does not leverage non-strict evaluation. We also note the non-strict language Clean [Brus et al. 1987] that employs uniqueness types to ensure resources are used at most once.

Inspirations for Linear Core. Linear Mini-Core [Bernardy et al. 2020] is a specification of linear types in Core as they were being implemented in GHC, and doubles as the (unpublished) precursor to our work. Linear Mini-Core first observes the incapacity of Core's type system to accept linear programs after transformations, and introduces usage environments for let-bound variables with the same goal of Linear Core of specifying a linear type system for Core that accepts the optimising transformations. We draw from Linear Mini-Core the rule for non-recursive let expressions and how let-bound variables are annotated with a usage environment. However, our work further explores the interaction of laziness with linearity in depth, and diverges significantly in rules for typing other constructs, notably, case expressions and case alternatives. Furthermore, unlike Mini-Core, we prove Linear Core is type safe, guarantees linear resource usage, and that multiple optimising transformations, when applied to Linear Core programs, preserve linearity. Our irrelevant resources have similarities with 0-multiplicity values from QTT [Brady 2021] which are disallowed from being used directly in executable code and, as previously noted, our tagged resources are related to fractional permissions from separation logic [Boyland 2003].

Linearity-directed optimisations. Core-to-Core transformations appear in multiple papers [Baker-Finch et al. 2004; Breitner 2016; Maurer et al. 2017; Peyton Jones and Marlow 2002; Peyton Jones and Partain 1996; Peyton Jones and Santos 1997; Santos and Peyton Jones 1995; Sergey et al. 2017], all designed in the context of a typed language (Core) which does not have linear types. However, some authors [Bernardy et al. 2017b; Peyton Jones and Partain 1996; Peyton Jones and Santos 1997] have observed that optimisations (in particular, let-floating and inlining) can greatly benefit from linearity analysis and, in order to improve those transformations, special purpose linear-type-inspired systems were created and implemented. Preserving linear types in Core throughout the compilation pipeline allows the optimiser to leverage non-heuristic linearity information where, previously, it would resort solely to ad-hoc or incomplete custom-built linearity inference passes (naturally, these passes are still necessary to optimise programs not using explicit linearity or multiplicity annotations). Linearity may potentially be used to the benefit of other transformations. An obvious candidate is *inlining*, which is applied based on heuristics from information provided by the *cardinality analysis* pass that counts occurrences of bound variables. Linearity can be used to non-heuristically inform the inliner [Bernardy et al. 2017b].

6.1 Future Work

We highlight some avenues of future work. Briefly, these include *multiplicity coercions*, optimisations leveraging linearity, and incorporating Linear Core in GHC Haskell.

Multiplicity Coercions. Linear Core does not have type equality coercions, a flagship feature of GHC Core's type system. Coercions allow the Core intermediate language to encode a panoply of Haskell source type-level features such as GADTs, type families, or newtypes. In Linear Haskell, multiplicities are introduced as annotations to function arrows which specify the linearity of the function. In practice, multiplicities are simply types of kind Multiplicity, where One and Many are the type constructors of the kind Multiplicity; multiplicity polymorphism follows from type polymorphism, where multiplicity variables are just type variables. Encoding multiplicities as types allows Haskell programs to leverage type-level features when handling multiplicities. For instance, it is possible via the use of GADTs to define a function whose linearity of its second argument depends on the value of its first argument, internally realized through so-called multiplicity coercions. Currently, Core cannot make use of such coercions to determine whether the usages of linear resources match their intended multiplicity. Studying the interaction between coercions and multiplicities is a main avenue of future work for Linear Core.

Linear Core in the Glasgow Haskell Compiler. Integrating Linear Core in the Glasgow Haskell Compiler is one of the ultimate goals of our work. Core's current type system ignores linearity due to its limitation in understanding semantic linearity, and our work fills this gap and would allow Core to be linearly typed all throughout. Implementing Linear Core in GHC is a challenging endeavour, since we must account for all other Core features (e.g. strict constructor fields), propagate new or modified data throughout the entire pipeline, and accept more optimisations. Despite our initiative in this direction³, we leave this as future work.

Generalizing Linear Core to Haskell. Linear types, despite their compile-time correctness guarantees regarding resource management, impose a burden on programmers in being a restrictive typing discipline (witnessed, e.g., by Linear Constraints [Spiwack et al. 2022]). Linear Core eases the restrictions of linear typing by being more flexible in understanding linearity for lazy languages. It is future work to apply our ideas to the surface Haskell language.

Acknowledgments

We would like to thank Arnaud Spiwack for many insightful discussions on linearity in Core and the reviewers for their feedback, which improved the paper significantly. This work was supported by Well-Typed LLP and by national funds through Fundação para a Ciência e a Tecnologia, I.P. (FCT) under projects UID/50021/2025 and UID/PRR/50021/2025.

References

Jean-Marc Andreoli. 1992. Logic Programming with Focusing Proofs in Linear Logic. Journal of Logic and Computation 2, 3 (06 1992), 297–347. arXiv:https://academic.oup.com/logcom/article-pdf/2/3/297/6137548/2-3-297.pdf doi:10.1093/logcom/2.3.297

Clem Baker-Finch, Kevin Glynn, and Simon Peyton Jones. 2004. Constructed Product Result Analysis for Haskell. *Journal of Functional Programming* 14, 2 (March 2004), 211–245. https://www.microsoft.com/en-us/research/publication/constructed-product-result-analysis-haskell/

Andrew Barber. 1996. *Dual Intuitionistic Linear Logic.* Technical Report ECS-LFCS-96-347. The University of Edinburgh. Jean-Philippe Bernardy, Mathieu Boespflug, Ryan R. Newton, Simon Peyton Jones, and Arnaud Spiwack. 2017a. Linear Haskell: practical linearity in a higher-order polymorphic language. *CoRR* abs/1710.09756 (2017). arXiv:1710.09756 http://arxiv.org/abs/1710.09756

³https://gitlab.haskell.org/ghc/ghc/-/merge requests/10310

- J. Bernardy, R. Eisenberg, M. Boespflug, R. Newton, S. Peyton Jones, and A. Spiwack. 2020. Linear Mini-Core. https://gitlab.haskell.org/ghc/ghc/-/wikis/uploads/355cd9a03291a852a518b0cb42f960b4/minicore.pdf.
- Jean-Philippe Bernardy, Mathieu Boespflug, Ryan R. Newton, Simon Peyton Jones, and Arnaud Spiwack. 2017b. Linear Haskell: practical linearity in a higher-order polymorphic language. Proc. ACM Program. Lang. 2, POPL, Article 5 (Dec. 2017), 29 pages. doi:10.1145/3158093
- Jean-Philippe Bernardy and Arnaud Spiwack. 2021. Evaluating Linear Functions to Symmetric Monoidal Categories. In *Proceedings of the 14th ACM SIGPLAN International Symposium on Haskell.* Association for Computing Machinery, New York, NY, USA, 14–26. doi:10.1145/3471874.3472980
- John Boyland. 2003. Checking interference with fractional permissions. In *Proceedings of the 10th International Conference on Static Analysis* (San Diego, CA, USA) (SAS'03). Springer-Verlag, Berlin, Heidelberg, 55–72.
- Edwin Brady. 2021. Idris 2: Quantitative Type Theory in Practice. In 35th European Conference on Object-Oriented Programming (ECOOP 2021) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 194), Anders Møller and Manu Sridharan (Eds.). Schloss Dagstuhl Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 9:1–9:26. doi:10.4230/LIPIcs.ECOOP.2021.9
- Joachim Breitner. 2016. Lazy Evaluation: From natural semantics to a machine-checked compiler transformation. Ph. D. Dissertation. Karlsruher Institut für Technologie (KIT). doi:10.5445/IR/1000054251
- T. H. Brus, M. C. J. D. van Eekelen, M. O. van Leer, and M. J. Plasmeijer. 1987. Clean A language for functional graph rewriting. In *Functional Programming Languages and Computer Architecture*, Gilles Kahn (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 364–384.
- Luís Caires and Frank Pfenning. 2010. Session Types as Intuitionistic Linear Propositions. In CONCUR 2010 Concurrency Theory, Paul Gastin and François Laroussinie (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 222–236.
- Iliano Cervesato, Joshua S. Hodas, and Frank Pfenning. 2000. Efficient resource management for linear logic proof search. Theoretical Computer Science 232, 1 (2000), 133–163. doi:10.1016/S0304-3975(99)00173-5
- Jawahar Chirimar, Carl A. Gunter, and Jon G. Riecke. 1996. Reference counting as a computational interpretation of linear logic. Journal of Functional Programming 6, 2 (1996), 195–244. doi:10.1017/S0956796800001660
- Pritam Choudhury, Harley Eades III, Richard A. Eisenberg, and Stephanie Weirich. 2021. A graded dependent type system with a usage-aware semantics. *Proc. ACM Program. Lang.* 5, POPL (2021), 1–32. doi:10.1145/3434331
- Duncan Coutts, Roman Leshchinskiy, and Don Stewart. 2007. Stream fusion: from lists to streams to nothing at all. In *Proceedings of the 12th ACM SIGPLAN International Conference on Functional Programming* (Freiburg, Germany) (*ICFP* '07). Association for Computing Machinery, New York, NY, USA, 315–326. doi:10.1145/1291151.1291199
- Luís Damas and Robin Milner. 1982. Principal Type-Schemes for Functional Programs. In Conference Record of the Ninth Annual ACM Symposium on Principles of Programming Languages, Albuquerque, New Mexico, USA, January 1982, Richard A. DeMillo (Ed.). ACM Press, 207–212. doi:10.1145/582153.582176
- Roy Dyckhoff. 1992. Contraction-Free Sequent Calculi for Intuitionistic Logic. *The Journal of Symbolic Logic* 57, 3 (1992), 795–807. http://www.jstor.org/stable/2275431
- Peng Fu, Kohei Kishida, and Peter Selinger. 2020. Linear Dependent Type Theory for Quantum Programming Languages: Extended Abstract. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science* (Saarbrücken, Germany) (*LICS '20*). Association for Computing Machinery, New York, NY, USA, 440–453. doi:10.1145/3373718.3394765
- Andrew Gill, John Launchbury, and Simon L. Peyton Jones. 1993. A short cut to deforestation. In *Proceedings of the Conference on Functional Programming Languages and Computer Architecture* (Copenhagen, Denmark) (*FPCA '93*). Association for Computing Machinery, New York, NY, USA, 223–232. doi:10.1145/165180.165214
- Jean-Yves Girard. 1987. Linear logic. *Theoretical Computer Science* 50, 1 (1987), 1–101. doi:10.1016/0304-3975(87)90045-4 Carl A. Gunter and Didier Rémy. 1993. A Proof-Theoretic Assessment of Runtime Type Errors. https://api.semanticscholar.org/CorpusID:2871566 AT&T Bell Laboratories Technical Memo 11261-921230-43TM.
- Wen Kokke and Ornela Dardha. 2021. Deadlock-Free Session Types in Linear Haskell. In *Proceedings of the 14th ACM SIGPLAN International Symposium on Haskell* (Virtual, Republic of Korea) (*Haskell 2021*). Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3471874.3472979
- John Launchbury. 1993. A natural semantics for lazy evaluation. In Proceedings of the 20th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (Charleston, South Carolina, USA) (POPL '93). Association for Computing Machinery, New York, NY, USA, 144–154. doi:10.1145/158511.158618
- Andrew Lelechenko. 2024. text-builder-linear: Builder for Text and ByteString based on linear types. https://hackage.haskell.org/package/text-builder-linear
- Jialin Li, Andrea Lattuada, Yi Zhou, Jonathan Cameron, Jon Howell, Bryan Parno, and Chris Hawblitzel. 2022. Linear Types for Large-Scale Systems Verification. Proc. ACM Program. Lang. 6, OOPSLA1, Article 69 (apr 2022), 28 pages. doi:10.1145/3527313
- Danielle Marshall, Michael Vollmer, and Dominic Orchard. 2022. Linearity and Uniqueness: An Entente Cordiale. In Programming Languages and Systems 31st European Symposium on Programming, ESOP 2022, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Munich, Germany, April 2-7, 2022, Proceedings (Lecture

- Notes in Computer Science, Vol. 13240), Ilya Sergey (Ed.). Springer, 346-375. doi:10.1007/978-3-030-99336-8_13
- Nicholas D. Matsakis and Felix S. Klock. 2014. The Rust Language. Ada Lett. 34, 3 (oct 2014), 103–104. doi:10.1145/2692956.
- Luke Maurer, Zena Ariola, Paul Downen, and Simon Peyton Jones. 2017. Compiling without continuations. In *ACM Conference on Programming Languages Design and Implementation (PLDI'17)*. ACM, 482–494. https://www.microsoft.com/enus/research/publication/compiling-without-continuations/
- Rodrigo Mesquita and Bernardo Toninho. 2025. Linear Core: Theory and a GHC Plugin. https://hackage.haskell.org/package/linear-core-prototype-0.1.0.0
- Simon Peyton Jones and Simon Marlow. 2002. Secrets of the Glasgow Haskell Compiler inliner. *Journal of Functional Programming* 12 (July 2002), 393–434. https://www.microsoft.com/en-us/research/publication/secrets-of-the-glasgow-haskell-compiler-inliner/
- Simon Peyton Jones and Will Partain. 1996. Let-Floating: Moving Bindings to Give Faster Programs. *Proc. of ICFP'96* 31 (10 1996). doi:10.1145/232629.232630
- Simon Peyton Jones and Andre Santos. 1997. A transformation-based optimiser for Haskell. *Science of Computer Programming* 32, 1 (October 1997). https://www.microsoft.com/en-us/research/publication/a-transformation-based-optimiser-for-haskell/
- Simon L. Peyton Jones. 1987. The Implementation of Functional Programming Languages (Prentice-Hall International Series in Computer Science). Prentice-Hall, Inc., USA.
- Matthew Pickering, Nicolas Wu, and Boldizsár Németh. 2019. Working with Source Plugins. In *Proceedings of the 12th ACM SIGPLAN International Symposium on Haskell* (Berlin, Germany) (*Haskell 2019*). Association for Computing Machinery, New York, NY, USA, 85–97. doi:10.1145/3331545.3342599
- Benjamin C. Pierce. 2004. Advanced Topics in Types and Programming Languages. The MIT Press.
- Andre Santos and Simon Peyton Jones. 1995. Compilation by transformation for non-strict functional languages. Ph. D. Dissertation. https://www.microsoft.com/en-us/research/publication/compilation-transformation-non-strict-functional-languages/
- Ilya Sergey, Dimitrios Vytiniotis, Simon L. Peyton Jones, and Joachim Breitner. 2017. Modular, higher order cardinality analysis in theory and practice. *Journal of Functional Programming* 27 (2017), e11. doi:10.1017/S0956796817000016
- Arnaud Spiwack, David Feuer, and José Pedro Magalhães. 2025. linear-generics: Generic programming library for generalised deriving. https://hackage.haskell.org/package/linear-generics
- Arnaud Spiwack, Csongor Kiss, Jean-Philippe Bernardy, Nicolas Wu, and Richard A. Eisenberg. 2022. Linearly Qualified Types: Generic Inference for Capabilities and Uniqueness. *Proc. ACM Program. Lang.* 6, ICFP, Article 95 (aug 2022), 28 pages. doi:10.1145/3547626
- Martin Sulzmann, Manuel M. T. Chakravarty, Simon Peyton Jones, and Kevin Donnelly. 2007. System F with type equality coercions. In *Proceedings of the 2007 ACM SIGPLAN International Workshop on Types in Languages Design and Implementation* (Nice, Nice, France) (TLDI '07). Association for Computing Machinery, New York, NY, USA, 53–66. doi:10.1145/1190315.1190324
- David N. Turner and Philip Wadler. 1999. Operational interpretations of linear logic. *Theoretical Computer Science* 227, 1 (1999), 231–248. doi:10.1016/S0304-3975(99)00054-7
- Philip Wadler. 1990. Linear Types can Change the World!. In Programming Concepts and Methods.

```
x^*
(\Lambda p. e)^*
                                                                       \triangleq \Lambda p. (e)^*
                                                                       \triangleq \lambda x :_{\pi} \sigma.(e)^*
(\lambda x :_{\pi} \sigma.e)^*
(e \pi)^*
                                                                       \triangleq (e)^* \pi
(e x)^*
                                                                        \triangleq (e)^* x
                                                                        \triangleq let x:_{\Lambda} \sigma = (e')^* in (e)^* x
(e \ e')^*
                                                                                              with \Gamma; \Delta' \vdash e : \sigma \rightarrow_{\pi} \tau
                                                                                              and \Gamma; \Delta \vdash e' : \sigma
                                                                       \triangleq let \overline{x_i :_{\Delta_i} \sigma_i = (e_i)^*} in K \overline{x_i}
K \overline{e_i}
                                                                                             with \overline{\Gamma; \Delta_i \vdash e_i : \sigma_i}
(case e of z:_{\Delta}\sigma \{\overline{\rho \Rightarrow e'}\})* \triangleq case (e)^* of z:_{\Delta}\sigma \{\overline{\rho \Rightarrow (e')^*}\}
(\operatorname{let} x:_{\Lambda} \sigma = e \text{ in } e')^* \qquad \qquad \triangleq \quad \operatorname{let} x:_{\Lambda} \sigma = (e)^* \text{ in } (e')^*
(\text{let rec } \overline{x:_{\Lambda} \sigma = e} \text{ in } e')^* \qquad \triangleq \quad \text{let rec } \overline{x:_{\Lambda} \sigma = (e)^*} \text{ in } (e')^*
```

Fig. 10. Explicit Sharing Translation

$$\frac{\Theta: e \Downarrow \Theta': \Lambda p.e' \quad \Theta': e'[\pi/p] \Downarrow \Theta'': v}{\Theta: e \pi \Downarrow \Theta': v}$$

$$\frac{\Theta: e \Downarrow \Theta': \lambda y:_{\pi} \sigma.e' \quad \Theta': e'[x/y] \Downarrow \Theta'': v}{\Theta: e \pi \Downarrow \Theta'': v}$$

$$\frac{(\Theta, x:_{\omega} \sigma = e): e \Downarrow \Theta': v}{(\Theta, x:_{\omega} \sigma = e): x \Downarrow (\Theta', x:_{\omega} \sigma = v): v} \qquad \frac{(\Theta, x:_{\omega} \sigma = e): e' \Downarrow \Theta': v}{\Theta: let x:_{\Delta} \sigma = e \text{ in } e' \Downarrow \Theta': v}$$

$$\frac{\Theta: e \Downarrow \Theta': K \overline{x_i} \quad \Theta': e'[\overline{x_i/y_i}] [K \overline{x_i/z}] \Downarrow \Theta'': v}{\Theta: case \ e \text{ of } z:_{\Delta'} \sigma \ \{\dots, K \overline{y_i} \rightarrow e', \dots\} \Downarrow \Theta'': v}$$

$$\frac{\Theta: e \Downarrow \Theta': K \overline{x_i} \quad \Theta': e'[K \overline{x_i/z}] \Downarrow \Theta'': v}{\Theta: case \ e \text{ of } z:_{\Delta'} \sigma \ \{\dots, K \overline{y_i} \rightarrow e', \dots\} \Downarrow \Theta'': v}$$

$$\frac{\Theta: e \Downarrow \Theta': K \overline{x_i} \quad \Theta': e'[K \overline{x_i/z}] \Downarrow \Theta'': v}{\Theta: case \ e \text{ of } z:_{\Delta'} \sigma \ \{\dots, K \overline{y_i} \rightarrow e', \dots\} \Downarrow \Theta'': v}$$

Fig. 11. Natural Semantics for λ_{Λ}^{π}

A Natural and Instrumented Operational Semantics

We define a lazy big-step operational semantics for λ_{Δ}^{π} in the style of Launchbury [Launchbury 1993] natural semantics. The semantics is equipped with an environment which assigns expressions to variables and is mutated in order to express lazy evaluation. Terms are explicit about their sharing (in the sense of shared reductions in lazy evaluation), and so we perform a translation that makes all sharing explicit through let-binders. This translation is presented in Figure 10. We write let $\overline{x_i}$: $\overline{A_i}$, $\overline{A_i}$ = $\overline{A_i}$ in e' for the iterated let-binding of variables $\overline{x_i}$.

The natural semantics (Figure 11) is defined via the relation $\Theta: e \Downarrow \Theta': v$ where e is an expression, v a value, Θ and Θ' are environments with bindings of the form $x:_{\omega} \sigma = e$ which assigns the expression e to the variable e of the given type.

Our instrumented, linearity-aware semantics (akin to that of [Bernardy et al. 2017a,b]) is defined in Figure 12. In this semantics, linear variables are erased from the runtime environment, ensuring

$$\overline{\Gamma; \Delta \vdash (\Theta \mid \Delta p.e) \Downarrow (\Theta \mid \Delta p.e) : \forall p.\sigma, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid \Delta p.e') : \forall p.\sigma, \Sigma} \qquad \Gamma; \Delta \vdash (\Theta' \mid e'[\pi/p]) \Downarrow (\Theta'' \mid v) : \sigma[\pi/p], \Sigma}$$

$$\overline{\Gamma; \Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid \lambda w :_{\pi} \sigma.e) \Downarrow (\Theta \mid \lambda x :_{\pi} \sigma.e) : \sigma \to_{\pi} \tau, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid \lambda y :_{\sigma} c.e') : \sigma \to_{1} \tau, x : \sigma, \Sigma} \qquad \Gamma; \Delta \vdash (\Theta', y :_{1} \sigma = x \mid e') \Downarrow (\Theta'' \mid v) : \tau, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid \lambda y :_{\omega} \sigma.e') : \sigma \to_{\omega} \tau, x : \sigma, \Sigma} \qquad \Gamma; \Delta \vdash (\Theta' \mid e'[x/y]) \Downarrow (\Theta'' \mid v) : \tau, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid \lambda y :_{\omega} \sigma.e') : \sigma \to_{\omega} \tau, x : \sigma, \Sigma} \qquad \Gamma; \Delta \vdash (\Theta' \mid e'[x/y]) \Downarrow (\Theta'' \mid v) : \tau, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid \lambda y :_{\omega} \sigma.e') : \sigma \to_{\omega} \tau, x : \sigma, \Sigma} \qquad \Gamma; \Delta \vdash (\Theta' \mid e'[x/y]) \Downarrow (\Theta'' \mid v) : \tau, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid v) : \sigma, \Sigma} \qquad \Delta'' = \Delta' \setminus (\Theta \setminus \Theta')$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid x :_{\Delta'} \sigma = e \mid x) \Downarrow (\Theta' \mid v) : \sigma, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid x :_{\Delta'} \sigma = e \mid x) \Downarrow (\Theta' \mid v) : \sigma, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid k x :_{\Delta'} \sigma = e \mid x) \Downarrow (\Theta' \mid v) : \tau, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid k x :_{\Delta'} \sigma = e \mid e') \Downarrow (\Theta' \mid v) : \tau, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid k x :_{\Delta'} \sigma = e \mid e') \Downarrow (\Theta' \mid v) : \tau, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid k x :_{\Delta'} \sigma = e \mid e') \Downarrow (\Theta' \mid v) : \tau, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid case \ e \ of \ z:_{\Delta'} T \ \{..., K \ y_{i,j} \implies e', ...\}\} \Downarrow (\Theta'' \mid v) : \tau, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid case \ e \ of \ z:_{\Delta'} T \ \{..., K \ y_{i,j} \implies e', ...\}\} \Downarrow (\Theta'' \mid v) : \tau, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid K x_{i,j}) : \tau, e', \Sigma} \quad \Gamma; \Delta \vdash (\Theta \mid e' \mid k x_{i,j} z_{i,j}) \Downarrow (\Theta'' \mid v) : \tau, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid K x_{i,j}) : \tau, e', \Sigma} \quad \Gamma; \Delta \vdash (\Theta \mid e' \mid E' x_{i,j} z_{i,j}) \Downarrow (\Theta'' \mid v) : \tau, \Sigma}$$

$$\underline{\Gamma; \Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid K x_{i,j}) : \tau, e', \Sigma} \quad \Gamma; \Delta \vdash (\Theta \mid e' \mid E' x_{i,j} z_{i,j}) \Downarrow (\Theta'' \mid v) : \tau$$

Fig. 12. Instrumented Semantics for λ_{Λ}^{π}

that any term violating linearity produces a stuck state. The semantics is defined by the judgment Γ ; $\Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid v) : \sigma$, Σ where Γ and Δ are typing judgments (Γ tracks unrestricted variables, Δ tracks linear assumptions) for variables occurring in Σ ; Θ and Θ' are runtime environments, consisting of bindings of the form $x :_{\pi} \sigma = e$ or $x :_{\Delta} \sigma = e$; e is the expression being evaluated and v its resulting value, both of type σ ; Σ is a list of assignments of the form $e : \tau$, which are expressions in which variables in Θ are used.

Following [Bernardy et al. 2017a; Gunter and Rémy 1993], we also consider partial derivations of our semantics in order to be able to state a progress result. Given our set of rules defining judgments of the form $a \downarrow b$ with ordered premises, we define a total derivation of $a \downarrow b$ as a tree in the standard way. A partial derivation $a \downarrow l$? is either a root labelled with $a \downarrow l$? or an application of a rule matching a on the left of the evaluation arrow, where exactly one of the premises $a' \downarrow l$? has a partial derivation, all the premises to the left of $a' \downarrow l$? have a total derivation, and the premises to the right of $a' \downarrow l$? are not known yet.

By construction, in a partial derivation, there is only one premise $b \Downarrow ?$ with no sub-derivation. We call b the head of such a partial derivation. We write $a \Downarrow^* b$ for the relation which holds when b is the head of some partial derivation with root $a \Downarrow ?$. We write $a \Downarrow^* b$ the partial evaluation relation.

Below, we write $\Theta \uparrow^1$ for the linear bindings in evaluation environment Θ and $\Theta \uparrow^{\omega}$ for all other bindings. $\Theta \downarrow_e^1$ refers to the expressions bound to linear variables in the runtime environment and $\Theta \downarrow_{\sigma}^1$ refers to the types of those expressions.

Definition A.1 (Environment Expansion). We rely on an expansion of a runtime environment Θ for the evaluation of a term e into a (typed) term, written $[\Theta \mid e]$, defined inductively as follows:

$$\begin{array}{lll} [\cdot \mid e] & \triangleq & e \\ [\Theta, x:_1 \tau = e' \mid e] & \triangleq & [\Theta \mid (\lambda x:_1 \tau.e) \ e'] \\ [\Theta, x:_{\Delta} \sigma = e' \mid e] & \triangleq & [\Theta \mid \text{let } x:_{\Delta} \sigma = e' \text{ in } e] \end{array}$$

We note that for the treatment of recursive let groups, environment expansion needs to happen for the entire group of let-bound variables. This is easily achieved by extending the runtime environment with information about recursive let groupings, which we omit for the sake of presentation.

Definition A.2 (Well-typed State). We write Γ ; $\Delta \vdash \Theta \mid e : \tau, \Sigma$ to denote the following typing judgment:

$$\Gamma; \Delta \vdash [\Theta \mid (e, terms(\Sigma))] : \tau \otimes \bigotimes(\Sigma)$$

where $terms(e_1:\sigma_1,\ldots,e_n:\sigma_n)$ stands for the tuple $(e_1,(\ldots,(e_n,())))$ and $\bigotimes(e_1:\sigma_1,\ldots,e_n:\sigma_n)$ for $\sigma_1\otimes(\ldots(\sigma_n\otimes()))$.

Lemma A.3.
$$\Gamma; \Delta \vdash [\Theta \mid e] : \sigma \text{ iff } \Gamma, \Theta \uparrow^{\omega}; \Delta, \Theta \uparrow^{1} \vdash (e, \Theta \downarrow^{1}_{e}) : \sigma \otimes \bigotimes (\Theta \downarrow^{1}_{\sigma})$$

Proof. By structural induction on the environment expansion definition in the forward direction and by a straightforward induction on Θ in the converse direction. We show only the forward direction.

```
Case: [\Theta, x:_1 \tau = y \mid e] \triangleq [\Theta \mid (\lambda x:_1 \tau.e) \mid y]
      (1) \Gamma; \Delta \vdash [\Theta, x:_1\tau = y \mid e] : \phi
      (2) \Gamma; \Delta \vdash [\Theta \mid (\lambda x:_1 \tau.e) \ y] : \phi
                                                                                                                                                                                                                                    by def.
      (3) \; \Gamma, \Theta \uparrow^\omega; \Delta, \Theta \uparrow^1 \vdash ((\lambda x :_1 \tau.e) \; y, \Theta \downarrow^1_e) : \phi \otimes \bigotimes (\Theta \downarrow^1_\sigma)
                                                                                                                                                                                                                                     by i.h.
      (4) Let \Delta_1, \Delta_2, \Delta_3 = \Delta, \Theta \uparrow^1
      (5) \Gamma, \Theta \uparrow^{\omega}; \Delta_3 \vdash \Theta \downarrow_e^1: \bigotimes (\Theta \downarrow_{\sigma}^1)
      (6) \Gamma, \Theta \uparrow^{\omega}; \Delta_2 \vdash y : \sigma
                                                                                                                                                                                                                                     by inv
      (7) \Gamma, \Theta \uparrow^{\omega}; \Delta_1 \vdash (\lambda x:_1 \tau.e) : \sigma \rightarrow \phi
                                                                                                                                                                                                                                     by inv
      (8) \Gamma, \Theta \uparrow^{\omega}; \Delta_1, x:_1 \tau \vdash e : \phi
                                                                                                                                                                                                                                     by inv
      (9) \Gamma, \Theta \uparrow^{\omega}; \Delta_1, \Delta_2, x:_1 \tau \vdash (e, y) : \phi \otimes \sigma
                                                                                                                                                                                                                                   by pair
      (10) \Gamma, \Theta \uparrow^{\omega}; \Delta, x:_1 \tau \vdash (e, y, \Theta \downarrow^1_e) : \phi \otimes \sigma \otimes \bigotimes (\Theta \downarrow^1_\sigma)
                                                                                                                                                                                                                                    by pair
```

A.1 Bisimulation

We relate our two operational semantics through a bisimulation, allowing us to transfer type safety of the instrumented semantics to the lazy big-step semantics. The relation essentially allows the execution environment to treat some unrestricted bindings of the natural semantics as linear.

Definition A.4 (Bisimulation). We define the relation $\gamma(\Theta:e)(\Gamma; \Delta \vdash \Theta' \mid e':\tau, \Sigma)$ between a well-typed state and an ordinary state of the lazy big-step semantics if e = e' and $\Theta' \uparrow^{\omega} \leq \Theta$.

LEMMA A.5.

- For all $\gamma(\Theta:e)(\Gamma; \Delta \vdash \Theta' \mid e:\tau, \Sigma)$ such that $\Theta:e \Downarrow \Theta'':v$ there exists a well-typed state $\Gamma; \Delta \vdash (\Theta''' \mid v):\tau, \Sigma$ such that $\Gamma; \Delta \vdash (\Theta' \mid e) \Downarrow (\Theta''' \mid v):\tau, \Sigma$ and $\gamma(\Theta'':v)(\Gamma; \Delta \vdash \Theta''' \mid v:\tau, \Sigma)$
- For all $\gamma(\Theta:e)(\Gamma; \Delta \vdash \Theta' \mid e:\tau, \Sigma)$ such that $\Theta:e \Downarrow^* \Theta'':v$ there exists a well-typed state $\Gamma; \Delta \vdash (\Theta''' \mid v):\tau, \Sigma$ such that $\Gamma; \Delta \vdash (\Theta' \mid e) \Downarrow^* (\Theta''' \mid v):\tau, \Sigma$ and $\gamma(\Theta'':v)(\Gamma; \Delta \vdash \Theta''' \mid v:\tau, \Sigma)$

PROOF. By induction on the given operational semantics rules.

LEMMA A.6.

- For all $\gamma(\Theta:e)(\Gamma; \Delta \vdash \Theta' \mid e:\tau, \Sigma)$ such that $\Gamma; \Delta \vdash (\Theta' \mid e) \Downarrow (\Theta'' \mid v):\tau, \Sigma$, there exists a state $\Theta''':v$ such that $\Theta:e \Downarrow \Theta''':v$ and $\gamma(\Theta''':v)(\Gamma; \Delta \vdash (\Theta'' \mid v):\tau, \Sigma)$
- For all $\gamma(\Theta:e)(\Gamma; \Delta \vdash \Theta' \mid e:\tau, \Sigma)$ such that $\Gamma; \Delta \vdash (\Theta' \mid e) \downarrow^* (\Theta'' \mid v):\tau, \Sigma$, there exists a state $\Theta''':v$ such that $\Theta:e \downarrow^* \Theta''':v$ and $\gamma(\Theta''':v)(\Gamma; \Delta \vdash (\Theta'' \mid v):\tau, \Sigma)$

PROOF. By induction on the given operational semantics rules.

B Type Safety Proofs

In the following we write \Rightarrow to stand for \Rightarrow_{WHNF} and \Rightarrow for \Rightarrow_{NWHNF} .

B.1 Auxiliary Lemmas

The following lemmas presuppose the variable names across the typing contexts are distinct.

Lemma B.1 (Linear to Δ -bound).

- (1) If Γ ; Δ' , $x:_1\tau \vdash e : \varphi$ then $\Gamma[\Delta/x]$, $x:_\Delta\sigma$; Δ , $\Delta' \vdash e : \varphi$ (with Δ fresh).
- (2) If Γ ; Δ' , x: $_1\tau \vdash_{alt} \rho \to e :_{\Delta_s}^z \sigma \Rightarrow \rho$ then $\Gamma[\Delta/x]$, x: $_\Delta\sigma$; Δ , $\Delta' \vdash_{alt} \rho \to e :_{\Delta_s}^z \sigma \Rightarrow \rho$ (with Δ fresh).

PROOF. Straightforward induction on the structure of the given derivation, reconstructing the original derivation and applying rule Var_{Δ} where rule Var_{1} was applied on x in the original derivation.

LEMMA 4.1 (LINEAR TO Δ-BOUND). If Γ ; Δ' , x:₁ $\tau \vdash e : \sigma$ then $\Gamma[\Delta/x]$, x: $\Delta \sigma$; Δ , $\Delta' \vdash e : \sigma$ (Δ fresh).

Lemma B.2 (Δ -bound to Linear).

- (1) If Γ , $x:_{\lceil \Delta \rceil} \tau$; $\lceil \Delta \rceil$, $\Delta' \vdash e : \varphi$ then $\Gamma[x/[\Delta]]$; Δ' , $x:_1 \tau \vdash e : \varphi$.
- (2) If $\Gamma, x:_{[\Delta]} \tau$; $[\Delta], \Delta' \vdash_{alt} \rho \to e :_{\Delta_s}^z \sigma \Rightarrow \rho \ then \Gamma[x/[\Delta]]; \Delta', x:_1 \tau \vdash_{alt} \rho \to e :_{\Delta_s}^z \sigma \Rightarrow \rho$.

PROOF. Straightforward induction on the structure of the given derivation, reconstructing the original derivation and applying rule Var_1 where rule Var_Δ was applied on x in the original derivation.

Lemma 4.2 (Δ-bound to Linear). If Γ , $x:_{[\Delta]}\tau$; $[\Delta]$, $\Delta' \vdash e : \sigma$ then $\Gamma[x/[\Delta]]$; Δ' , $x:_1\tau \vdash e : \sigma$.

Lemma B.3 (Unrestricted and Δ -bound).

- (1) $\Gamma, x:_{\omega}\tau; \Delta \vdash e : \varphi \text{ iff } \Gamma, x:_{\tau}\tau; \Delta \vdash e : \varphi$
- (2) $\Gamma, x:_{\omega}\tau; \Delta \vdash_{alt} \rho \rightarrow e:_{\Delta_s}^z \sigma \Rightarrow \rho \ iff \Gamma, x:\tau; \Delta \vdash_{alt} \rho \rightarrow e:_{\Delta_s}^z \sigma \Rightarrow \rho$

PROOF. Straightforward induction on the structure of the given derivation, noting that rules Var_{Δ} and Var_{ω} coincide when Δ is empty.

Lemma 4.3 (Unrestricted and Δ-bound). $\Gamma, x:_{\omega}\tau; \Delta \vdash e : \sigma \ iff \Gamma, x:_{\emptyset}\tau; \Delta \vdash e : \sigma$

B.2 Irrelevance

Lemma 4.4 (Irrelevance). If Γ , $z:_{[\Delta]}\tau$; $[\Delta]$, $\Delta' \vdash \rho \Rightarrow_{\mathsf{NWHNF}} e :_{[\Delta]}^z \sigma$ then Γ , $z:_{\Delta^\dagger}\tau$; Δ^\dagger , $\Delta' \vdash \rho \Rightarrow e :_{\Delta^\dagger}^z \sigma$, for any Δ^\dagger

PROOF. By structural induction on the case alternative typing derivation.

Case: Alt_

- $(1) \; \Gamma, z:_{[\Delta]} \tau; [\Delta] \; , \Delta' \vdash _ \Rightarrow e :_{[\Delta]}^z \sigma$
- (2) Γ , $z:_{\lceil \Delta \rceil} \tau$; $[\Delta]$, $\Delta' \vdash e : \sigma$
- (3) $[\Delta]$ is used through z

since $[\Delta]$ can't otherwise be used

and is introduced in this alternative uniquely

(since we have multi-tier proof irrelevance, i.e. $[[\Delta]] \neq [\Delta]$)

(4)
$$\Gamma[z/[\Delta]]$$
; $z:_1\tau, \Delta' \vdash e : \sigma$

by Lemma 4.2 (2,3)

(5)
$$\Gamma[\overline{\Delta_i}/[\Delta]]$$
, $z:_{\overline{\Delta_i}}\tau; \overline{\Delta_i}, \Delta' \vdash e : \sigma$

by Lemma 4.1 (4)

(6)
$$\Gamma$$
, $z:_{\overline{\Delta_i}}\tau$; $\overline{\Delta_i}$, $\Delta' \vdash e : \sigma$

by (3), resources in Γ with $[\Delta]$ (parts) in

the usage environment cannot be used in \boldsymbol{e}

(7)
$$\Gamma, z:_{\overline{\Delta_i}} \tau; \overline{\Delta_i}, \Delta' \vdash_{-} \implies e:_{\overline{\Delta_i}}^z \sigma$$

by *Alt*_

Case: Alt0

$$(1) \; \Gamma, z_{:[\Delta]} \tau; [\Delta] \; , \Delta' \vdash K \; \overline{x_{:\omega} \tau} \Rightarrow e \; :^{z}_{[\Lambda]} \; \sigma$$

(2)
$$\Gamma, z := \tau, \overline{x :_{\omega} \tau}; \Delta' \vdash e : \sigma$$

by inv. on Alt0 and def. of empty subst.

(3)
$$\Gamma, z:_{\overline{\Delta_i}} \tau; \overline{\Delta_i}, \Delta' \vdash K \overline{x:_{\omega} \tau} \Rightarrow e :_{\overline{\Delta_i}}^z \sigma$$

by *Alt*0 (2)

Case: AltN_{WHNF}

Not applicable since \Rightarrow is only generalized by \Rightarrow , not \Rightarrow .

Case: Alt N_{Not WHNF} We prove the theorem for constructing both an Alt N_{Not WHNF} and an $Alt N_{WHNF}$ from a proof-irrelevant $Alt N_{Not WHNF}$, to prove the statement holds for any \Rightarrow kind for Alt N, rather than just \Rightarrow or \Rightarrow .

$$(1)\;\Gamma,z:_{[\Delta]}\tau;[\Delta]\;,\Delta'\vdash K\;\overline{x:_{\omega}\tau},\overline{y_{i}:_{1}\sigma_{i}}^{n}\Rightarrow e:_{[\Delta]}^{z}\sigma$$

(2)
$$\Gamma, z:_{[\Delta]} \tau, \overline{x:_{\omega} \tau}, \overline{y_{j}:_{\Delta_{j}} \sigma_{j}}^{n}; [\Delta], \Delta' \vdash e : \sigma$$

by inv. on $Alt N_{\text{Not WHNF}}$

$$(3) \ \overline{\Delta_j} = \overline{[\Delta] \# K_j}$$

by inv. on Alt N_{Not WHNF}

Subcase
$$[\Delta]$$
 is consumed through z
(4) $\Gamma, \overline{x}:_{\omega}\tau, \overline{y}_{i}:_{\Delta,[z/[\Delta]]}\overline{\sigma}_{i}^{n}; z:_{1}\tau, \Delta' \vdash e : \sigma$

by Lemma 4.2 (2,subcase)

and vars in Γ with $[\Delta]$ cannot be used by subcase

Subcase constructing $AltN_{WHNF}$ (\Longrightarrow)

(5)
$$\Gamma$$
, $\overline{x}:_{\omega}\tau$, $\overline{y}_{j}:_{\Delta_{i}}\sigma_{j}^{n}$; $z:_{1}\tau$, $\Delta' \vdash e : \sigma$

by \overline{y} is not used, and Weaken

$$(6)\;\Gamma,\overline{x:_{\omega}\tau},\overline{y_{j}:_{\Delta_{i}}\sigma_{j}}^{n},z:_{\overline{\Delta_{i}}}\tau;\overline{\Delta_{i}},\Delta'\vdash e:\sigma$$

by Lemma 4.1 (5) and Γ vars do not mention z by (4)

$$(7) \; \Gamma, z:_{\overline{\Delta_i}} \tau; \overline{\Delta_i}, \Delta' \vdash K \; \overline{x:_{\omega} \tau}, \overline{y_i:_1 \sigma_i}^n \; \Longrightarrow \; e \; :_{\overline{\Delta_i}}^z \sigma$$

by $Alt N_{WHNF}$ (6)

by Lemma 4.1 (5)

Subcase constructing $AltN_{\text{Not WHNF}}$ (\Rightarrow)

(5)
$$\Gamma$$
, $\overline{x}:_{\omega}\tau$, $\overline{y}_{j}:_{\overline{\Delta_{i}}\#K_{j}}\sigma_{j}^{n}$; $z:_{1}\tau$, $\Delta' \vdash e:\sigma$

by \overline{y} does not consume resources

(6)
$$\Gamma, \overline{x}:_{\omega} \overline{\tau}, \overline{y}_{j}:_{\overline{\Delta_{i}}\#K_{i}} \overline{\sigma_{j}}^{n}, z:_{\overline{\Delta_{i}}} \tau; \overline{\Delta_{i}}, \Delta' \vdash e : \sigma$$

and Γ vars do not mention z by (4)

$$(7) \; \Gamma, z:_{\overline{\Delta_i}} \tau; \overline{\Delta_i}, \Delta' \vdash K \; \overline{x:_{\omega} \tau}, \overline{y_i:_1 \sigma_i}^n \Rightarrow e \; :_{\overline{\Delta_i}}^z \sigma$$

by $Alt N_{\text{Not WHNF}}$ (6)

Subcase $[\Delta]$ is (fully) consumed by \overline{y} (after splitting)

$$(4) \; \Gamma, z:_{\overline{y}}\tau, \overline{x:_{\omega}\tau}; \overline{y_{j}:_{1}\sigma_{j}}^{n}, \Delta' \vdash e: \sigma$$

by Lemma 4.2 (2, subcase)

and vars in Γ with $[\Delta]$ cannot be used by subcase by Weaken and z does not consume resources

(5)
$$\Gamma$$
, $z:_{\overline{\Delta_i}}\tau$, $\overline{x}:_{\omega}\tau$; $\overline{y}_j:_{1}\overline{\sigma_j}^n$, $\Delta' \vdash e : \sigma$
Subcase constructing $AltN_{WHNF}$ (\Longrightarrow)

by Lemma 4.1 (5)

(6)
$$\Gamma, z:_{\overline{\Delta_i}} \tau, \overline{x:_{\omega} \tau}, \overline{y_j:_{\Delta_i} \sigma_j}^n; \overline{\Delta_i}, \Delta' \vdash e : \sigma$$

and vars in Γ do not mention \overline{y} (4,5)

$$(7) \Gamma, z:_{\overline{\Delta_i}} \tau; \overline{\Delta_i}, \Delta' \vdash K \overline{x:_{\omega} \tau}, \overline{y_i:_1 \sigma_i}^n \Longrightarrow e :_{\overline{\Delta_i}}^z \sigma$$

by AltN_{WHNF}

Subcase constructing
$$Alt N_{\text{Not WHNF}} (\Longrightarrow)$$

(6) $\Gamma, z:_{\overline{\Delta_i}} \tau, \overline{x}:_{\omega} \tau, \overline{y}_j:_{\overline{\Delta_i} \# K_i} \sigma_j^n; \overline{\Delta_i}, \Delta' \vdash e : \sigma$

by $1 \Rightarrow \Delta$ lemma (5)

(7)
$$\Gamma, z:_{\overline{\Delta_i}} \tau; \overline{\Delta_i}, \Delta' \vdash K \ \overline{x:_{\omega} \tau}, \overline{y_i:_1 \sigma_i}^n \Rightarrow e :_{\overline{\Delta_i}}^z \sigma$$

and vars in Γ do not mention \overline{y} (4,5) by $Alt N_{\text{Not WHNF}}$

Substitution Lemmas

The linear substitution lemma states that a well-typed expression *e* with a linear variable *x* of type σ remains well-typed if occurrences of x in the e are replaced by an expression e' of the same type

 σ , and occurrences of x in the linear context and in usage environments of Δ -bound variables are replaced by the linear context Δ' used to type e':

LEMMA B.4 (SUBSTITUTION OF LINEAR VARIABLES PRESERVES TYPING).

- (1) If Γ ; Δ , x:₁ $\tau \vdash e : \varphi$ and Γ ; $\Delta' \vdash e : \sigma$ then $\Gamma [\Delta'/x] ; \Delta$, $\Delta' \vdash e[e'/x] : \varphi$
- (2) If $\Gamma; \Delta, x:_1 \tau \vdash \rho \implies e:_{\Delta_s}^z \sigma \text{ and } \Gamma; \Delta' \vdash e': \tau \text{ and } \Delta_s \subseteq \Delta, x \text{ then } \Gamma\left[\Delta'/x\right]; \Delta, \Delta' \vdash \rho \implies$ $e[e'/x]:^z_{\Lambda_c[\Lambda'/x]}\sigma$

Where $\Gamma[\Delta'/x]$ substitutes all occurrences of x in the usage environments of Δ -variables in Γ by the linear variables in Δ' . We further require that the environment annotated in the case alternative judgement, Δ_s , is a subset of the environment used to type the whole alternative $\Delta_s \subseteq \Delta$. In all occurrences of the alternative judgement (in CaseWHNF and CaseNot WHNF), the environment annotating the alternative judgement is *always* a subset of the alternative environment.

Proof. By structural induction on the given derivation. We omit most trivial cases. Statement (1):

Case: ∧I

```
(1) \Gamma; \Delta, x:<sub>1</sub>\sigma \vdash \Lambda p. e : \forall p. \varphi
```

(2) Γ ; $\Delta' \vdash e' : \sigma$

(3) Γ , p; Δ , x: $_1<math>\sigma \vdash e : \varphi$ by inversion on ΛI

(4) $p \notin \Gamma$ by inversion on ΛI

(5) $\Gamma[\Delta'/x], p; \Delta, \Delta' \vdash e[e'/x] : \varphi$ by i.h.(1) by (2,3)

(6) $\Gamma[\Delta'/x]$; $\Delta, \Delta' \vdash \Lambda p. e[e'/x] : \forall p. \varphi$ by ΛI (4,5)

(7) $(\Lambda p. e)[e'/x] = (\Lambda p. e[e'/x])$ by def. of substitution

Case: λI_1

(1)
$$\Gamma$$
; Δ , $x:_1\tau \vdash \lambda y:_1\sigma'$. $e:\sigma' \rightarrow_1 \varphi$

(2) Γ ; $\Delta' \vdash e' : \sigma$

(3)
$$\Gamma$$
; Δ , x :₁ τ , y :₁ $\sigma' \vdash e : \varphi$ by inversion on λI

(4) $\Gamma[\Delta'/x]$; Δ , y:₁ σ' , $\Delta' \vdash e[e'/x] : \varphi$ by i.h.(1) by (2,3)

(5) $\Gamma[\Delta'/x]$; $\Delta, \Delta' \vdash \lambda y$: $\sigma' \cdot e[e'/x] : \sigma' \rightarrow_1 \varphi$ by λI (4)

(6) $(\lambda y:_1\sigma'. e)[e'/x] = (\lambda y:_1\sigma'. e[e'/x])$ by def. of substitution

Case: Var₁

(1)
$$\Gamma$$
; x : $_1\sigma \vdash x : \sigma$

(2)
$$\Gamma$$
; $\Delta' \vdash e' : \sigma$

(3)
$$\Gamma[\Delta'/x]$$
; $\Delta' \vdash e' : \sigma$ by weaken

(4) x[e'/x] = e'by def. of substitution

(5) $\Gamma[\Delta'/x]$; $\Delta' \vdash e' : \sigma$ by (3)

Case: Var_{ω}

(1) Impossible. $x:_1\sigma$ can't be in the context.

Case: Var_{Λ}

(1)
$$\Gamma$$
, $y:_{\Delta,x:_1\tau}\varphi$; Δ , $x:_1\sigma \vdash y:\varphi$

(2)
$$\Gamma$$
; $\Delta' \vdash e' : \sigma$

(3)
$$y[e'/x] = y$$

by i.h.(1) (2,4)

(4)
$$\Gamma[\Delta'/x], y:_{\Delta,\Delta'}\varphi; \Delta, \Delta'; \vdash y:\varphi$$
 by Var_{Δ}

Case: λE_1

- (1) Γ ; Δ , Δ'' , x:₁ $\tau \vdash e \ e'' : \varphi$
- (2) $\Gamma: \Delta' \vdash e' : \sigma$

Subcase x occurs in e

- (3) Γ ; Δ , x:₁ $\tau \vdash e : \sigma' \rightarrow_1 \varphi$ by inversion on λE_1 (4) Γ ; $\Delta'' \vdash e'' : \sigma'$ by inversion on λE_1 by i.h.(1) (2,3)
- (5) $\Gamma[\Delta'/x]$; $\Delta, \Delta' \vdash e[e'/x] : \sigma' \rightarrow_1 \varphi$ (6) $\Gamma[\Delta'/x]$; $\Delta, \Delta', \Delta'' \vdash e[e'/x] e'' : \varphi$
- by λE_1 (7) (e[e'/x]e'') = (ee'')[e'/x]because x does not occur in e''

Subcase x occurs in e''

- (3) Γ ; $\Delta \vdash e : \sigma' \rightarrow_1 \varphi$ by inversion on λE_1 (4) Γ ; Δ'' , x: $_1\tau \vdash e'' : \sigma'$ by inversion on λE_1
- (5) $\Gamma[\Delta'/x]$; Δ'' , $\Delta' \vdash e''[e'/x] : \sigma'$ by i.h.(1) (2,4) (6) $\Gamma[\Delta'/x]$; Δ , Δ' , $\Delta'' \vdash e \ e''[e'/x] : \varphi$ by λE_1
- (7) (e e''[e'/x]) = (e e'')[e'/x]because x does not occur in e

Case: λE_{ω}

- (1) Γ ; Δ , x:₁ $\tau \vdash e e'' : \varphi$
- (2) Γ ; $\Delta' \vdash e' : \sigma$
- by e'' linear context is empty (3) x does not occur in e''
- (4) Γ ; Δ , $x:_1\tau \vdash e : \sigma' \rightarrow_{\omega} \varphi$ by inversion on λE_{ω} (5) Γ ; $\cdot \vdash e'' : \sigma'$ by inversion on λE_{ω}
- (6) $\Gamma[\Delta'/x]; \Delta, \Delta' \vdash e[e'/x] : \sigma' \rightarrow_{\omega} \varphi$
- (7) $\Gamma[\Delta'/x]$; $\Delta, \Delta' \vdash e[e'/x] e'' : \varphi$
- by λE_{ω} (8) (e[e'/x]e'') = (ee'')[e'/x]because x does not occur in e''

Case: Let

- (1) Γ ; $\Delta' \vdash e' : \sigma$
- Subcase x occurs in e
- (2) Γ ; Δ , x:₁ τ , Δ'' \vdash let y: $_{\Delta,x$:₁ $\tau}\sigma' = e$ in e'': φ
- (3) Γ , $y:_{\Delta,x:_1\tau}\sigma'$; Δ , $x:_1\tau$, $\Delta'' \vdash e'' : vp$ by inversion on Let
- (4) Γ ; Δ , $x:_1\tau \vdash e : \sigma'$ by inversion on Let
- (5) $\Gamma[\Delta'/x]$, $y:_{\Delta,\Delta'}\sigma'$; $\Delta,\Delta',\Delta'' \vdash e''[e'/x]$ by i.h.(1) (1, 3)
- (6) $\Gamma[\Delta'/x]$; $\Delta, \Delta' \vdash e[e'/x] : \sigma'$ by i.h.(1) (1, 4)
- (7) $\Gamma[\Delta'/x]$; $\Delta, \Delta', \Delta'' \vdash \text{let } y:_{\Delta,\Delta'} \sigma' = e[e'/x] \text{ in } e''[e'/x] : \varphi$ (5,6) by Let
- by subst.
- (8) (let $y:_{\Delta,\Delta'}\sigma' = e[e'/x]$ in e''[e'/x]) = (let $y:_{\Delta,\Delta'}\sigma' = e$ in e'')[e'/x]

Subcase x does not occur in e

- (2) Γ ; Δ , Δ'' , x:₁ $\tau \vdash$ **let** y: $\wedge \sigma' = e$ **in** $e'' : \varphi$
- (3) Γ , $y:_{\Lambda}\sigma'$; Δ , Δ'' , $x:_{1}\tau \vdash e'' : \varphi$ by inversion on Let (4) Γ ; $\Delta \vdash e : \sigma'$ by inversion on Let
- (5) $\Gamma[\Delta'/x], y: \Lambda \sigma'; \Delta, \Delta', \Delta'' \vdash e''[e'/x] : \varphi$ by i.h.(1) (1,3) (6) $\Gamma[\Delta'/x]$; $\Delta, \Delta', \Delta'' \vdash \text{let } y:_{\Delta} \sigma' = e \text{ in } e''[e'/x] : \varphi$ by Let (2,5,6)
- (7) let $y:\Delta\sigma'=e$ in $e''\lceil e'/x\rceil=(\text{let }y:\Delta\sigma'=e\text{ in }e'')\lceil e'/x\rceil$ by x does not occur in e

Case: LetRec

(1) Γ ; $\Delta' \vdash e' : \sigma$

Subcase $x:_1\tau$ occurs in some e_i

- (2) Γ ; Δ , x:₁ τ , Δ'' \vdash **let rec** $\overline{y_i:_{\Delta,x:_1\tau}\sigma_i=e_i}$ **in** $e'':\varphi$
- (3) Γ , $\overline{y_i:_{\Delta,x:_1\tau}\sigma_i}$; Δ , $x:_1\tau$, $\Delta'' \vdash e'' : \varphi$ by inversion on *LetRec* (4) Γ , $\overline{y_i:_{\Delta,x:_1\tau}\overline{\sigma_i}}$; Δ , $x:_1\tau \vdash e_i : \overline{\sigma_i}$
- by inversion on *LetRec* (5) $\Gamma[\Delta'/x], \overline{y_i:_{\Delta,\Delta'}\sigma_i}; \Delta, \Delta', \Delta'' \vdash e''[e'/x] : \varphi$ by i.h.(1) (1,3)
- (6) $\Gamma, \overline{y_i:_{\Delta,\Delta'}\sigma_i}; \Delta, \Delta' \vdash e_i[e'/x] : \sigma_i$ by i.h.(1) (1,4)
- (7) $\Gamma[\Delta'/x]$; $\Delta, \Delta', \Delta'' \vdash \text{let rec } \overline{y_{i:\Delta,\Gamma'_i}\sigma_i = e_i[e'/x]} \text{ in } e''[e'/x] : \varphi$ by LetRec
- (8) (let rec $\overline{y_i:_{\Delta,\Delta'}\sigma_i=e_i}$ in e'')[e'/x] = let rec $\overline{y_i:_{\Delta,\Delta'}\sigma_i=e_i[e'/x]}$ in e''[e'/x]

Subcase $x:_1\tau$ does not occur in any e_i

- (2) Γ ; Δ , $x:_1\tau$, $\Delta'' \vdash \text{let rec } \overline{y_i:_\Delta\sigma_i = e_i} \text{ in } e'' : \varphi$
- (3) Γ , $\overline{y_i:_{\Delta}\sigma_i}$; Δ , $x:_1\tau$, $\Delta'' \vdash e'' : \varphi$ by inversion on *LetRec* (4) $\overline{\Gamma, \overline{y_i}:_{\Delta}\sigma_i}; \Delta \vdash e_i : \sigma_i$
- by inversion on *LetRec* (5) $\Gamma[\Delta'/x], \overline{y_i:\Delta\sigma_i}; \Delta, \Delta', \Delta'' \vdash e''[e'/x] : \varphi$ by i.h.(1) (1,3)
- (6) $\Gamma[\Delta'/x]$; $\Delta, \Delta', \Delta'' \vdash \text{let rec } \overline{y_i:_\Delta \sigma_i = e_i} \text{ in } e''[e'/x] : \varphi$ by LetRec
- (7) let rec $\overline{y_i:_{\Delta}\sigma_i=e_i}$ in $e''[e'/x]=(\text{let rec }\overline{y_i:_{\Delta}\sigma_i=e_i}\text{ in }e'')[e'/x]$ by subcase

Case: CaseWHNF

(1) Γ ; $\Delta' \vdash e' : \sigma$

Subcase *x* occurs in *e*

- (2) Γ ; Δ , x:₁ τ , Δ'' \vdash **case** e **of** z: $_{\Delta,x$:₁ $\tau}\sigma'$ { $\rho \Rightarrow e''$ } : φ
- (3) e is in WHNF
- (4) Γ ; Δ , $x:_1\tau \vdash e : \sigma'$
- (5) ρ_{j} matches e (6) $\Gamma, z:_{\Delta,x:_{1}\tau}\sigma'; \Delta, x:_{1}\tau, \Delta'' \vdash \rho_{j} \Longrightarrow e'' :_{\Delta,x:_{1}\tau}^{z} \varphi$ (7) $\overline{\Gamma, z:_{[\Delta,x:_{1}\tau]}\sigma'; [\Delta, x:_{1}\tau], \Delta'' \vdash \rho \Longrightarrow e'' :_{[\Delta,x:_{1}\tau]}^{z} \varphi}$
- by inv. (8) $\Gamma \left[\Delta'/x \right]$; Δ , $\Delta' \vdash e[e'/x] : \varphi$
- by i.h.(1)
- (9) Γ $[\Delta'/x]$, $z:_{\Delta,\Delta'}\sigma';_{\Delta}\Delta',\Delta'' \vdash \rho_j \Rightarrow e''[e'/x]:_{\Delta,\Delta'}^z \varphi$ (10) Γ $[\Delta'/x]$, $z:_{[\Delta,\Delta']}\sigma';_{\Delta}\Delta']$, $\Delta'' \vdash_{alt} \rho \Rightarrow e''[e'/x]:_{[\Delta,\Delta']}^z \varphi$ by i.h.(2)
- by Irrelevance
- (11) $\Gamma \left[\Delta'/x \right]$; Δ , Δ' , $\Delta'' \vdash \mathbf{case} \ e[e'/x] \ \mathbf{of} \ z_{:\Delta,\Delta'}\sigma' \ \{ \overline{\rho \to e''[e'/x]} \} : \varphi$

Subcase *x* occurs in $\overline{e''}$

- (2) Γ ; Δ , Δ'' , $x:_1\tau \vdash \mathbf{case}\ e\ \mathbf{of}\ z:_\Delta\sigma'\ \{\overline{\rho \Rightarrow e''}\}: \varphi$
- (3) e is in WHNF
- (4) ρ_i matches e
- (5) Γ ; $\Delta \vdash e : \sigma'$
- (6) $\Gamma, z_{:\Delta}\sigma'; \Delta, \Delta'', x_{:1}\tau \vdash_{alt} \rho_j \to e'' :_{\Delta}^z \sigma' \Longrightarrow \varphi$ (7) $\Gamma, z_{:[\Delta]}\sigma'; [\Delta], \Delta'', x_{:1}\tau \vdash \rho \Longrightarrow e'' :_{[\Delta]}^z \varphi$ by inv.
- (8) e[e'/x] = eby x does not occur in e
- (9) $\Gamma[\Delta'/x], z:_{\Delta}\sigma'; \Delta, \Delta'', \Delta' + \rho_j \Rightarrow e''[e'/x] :_{\Delta}^{z} \varphi$ (10) $\Gamma[\Delta'/x], z:_{[\Delta]}\sigma'; [\Delta], \Delta'', \Delta' + \rho \Rightarrow e''[e'/x] :_{[\Delta]}^{z} \varphi$ by i.h.(2) by i.h.(2)
- (11) $\Gamma \left[\Delta'/x \right]$; $\Delta, \Delta'', \Delta' \vdash \mathbf{case} \ e \ \mathbf{of} \ z_{\Delta} \sigma' \ \{ \overline{\rho \Rightarrow e'' [e'/x]} \} : \varphi$

Case: CaseNotWHNF

(1) Γ ; $\Delta' \vdash e' : \sigma$

by inv.

Subcase *x* occurs in *e*

(2)
$$\Gamma$$
; Δ , x : $_1\tau$, Δ'' \vdash case e of z : $_{[\Delta,x:_1\tau]}\sigma'$ $\{\overline{\rho \Rightarrow e''}\}: \varphi$

(3) e is definitely not in WHNF

(4)
$$\Gamma$$
; Δ , x :₁ $\tau \vdash e : \sigma'$ by inv.

(5)
$$\overline{\Gamma, z_{:[\Delta, x_{:1}\tau]}} \sigma'; [\Delta, x_{:1}\tau], \Delta'' \vdash \rho \Rightarrow e'' :_{[\Delta, x_{:1}\tau]}^z \varphi$$
 by inv.

(6)
$$\Gamma[\Delta'/x]; \Delta, \Delta' \vdash e[e'/x] : \varphi$$
 by i.h.(1)

(7)
$$\Gamma[\Delta'/x], z_{:[\Delta,\Delta']}\sigma'; [\Delta,\Delta'], \Delta'' \vdash \rho \Rightarrow e''[e'/x] :_{[\Delta,\Delta']}^{z} \varphi$$
 by Irrelevance

(8)
$$\Gamma [\Delta'/x] ; \Delta, \Delta', \Delta'' \vdash \mathbf{case} \ e[e'/x] \ \mathbf{of} \ z_{:\Delta,\Delta'} \sigma' \ \{ \overline{\rho \Rightarrow e''[e'/x]} \} : \varphi$$

Subcase *x* occurs in $\overline{e''}$

(2)
$$\Gamma$$
; Δ , Δ'' , x : $_1\tau \vdash \mathbf{case}\ e\ \mathbf{of}\ z$: $_{[\Delta]}\sigma'\ \{\overline{\rho \Rightarrow e''}\}: \varphi$

(3) e is definitely not in WHNF

(4)
$$\Gamma$$
; $\Delta \vdash e : \sigma'$ by inv.

(5)
$$\Gamma, z:_{[\Delta]} \sigma'; [\Delta], \Delta'', x:_1 \tau \vdash \rho \Rightarrow e'' :_{[\Delta]}^z \varphi$$
 by inv.

(6)
$$e[e'/x] = e$$
 by x does not occur in e

(7)
$$\Gamma\left[\Delta'/x\right], z:_{[\Delta]}\sigma'; [\Delta], \Delta'', \Delta' \vdash \rho \Rightarrow e''[e'/x]:_{[\Delta]}^{z} \varphi$$
 by i.h.(2)

(8)
$$\Gamma \left[\Delta'/x \right]$$
; $\Delta, \Delta'', \Delta' \vdash \mathbf{case} \ e \ \mathbf{of} \ z_{\left[\Delta \right]} \sigma' \ \left\{ \overline{\rho \Rightarrow e''[e'/x]} \right\} : \varphi$

Statement (2):

Case: $AltN_{WHNF}$

(1)
$$\Gamma$$
; $\Delta' \vdash e' : \sigma$

(2)
$$\Gamma$$
; Δ , $x:_1\tau \vdash K \overline{x:_{\omega}\tau}$, $\overline{y_i:_1\sigma_i}^n \iff e:_{\Delta_e}^z \varphi[\dagger]$

(3)
$$\Gamma, \overline{x_{:\omega}\tau}, \overline{y_i:_{\Delta_i}\sigma_i}; \Delta, x_{:1}\tau \vdash e : \varphi$$
 by inv.

(4)
$$\Gamma\left[\Delta'/x\right], \overline{x_{:\omega}\tau}, \overline{y_{i:\Delta_i[\Delta'/x]}\sigma_i}; \Delta, \Delta' \vdash e[e'/x] : \varphi$$
 by i.h.(1)

(5)
$$\Gamma \left[\Delta'/x \right]$$
; $\Delta, \Delta' \vdash \rho \Longrightarrow e[e'/x] :_{\Delta_s[\Delta'/x]}^z \varphi[\dagger]$ by (4)

Case: AltN_{NotWHNF}

(1)
$$\Gamma$$
; $\Delta' \vdash e' : \sigma$

$$(2) \ \underline{\Gamma}; \underline{\Lambda}, \underline{x}_{:1} \underline{\tau} + K \ \overline{x}_{:\omega} \underline{\tau}, \underline{y}_{i:1} \underline{\sigma}_{i}^{n} \ \Rightarrow \ e :_{\underline{\Lambda}_{s}}^{z} \varphi[\ddagger]$$

$$(3) \overline{\Delta_i} = \overline{\Delta_s} \# K_j^n$$

(4) $\Gamma, \overline{x_{:\omega}\tau}, \overline{y_i:_{\Delta_i}\sigma_i}; \Delta, x:_1\tau \vdash e: \varphi$ by inv.

(5)
$$\Gamma\left[\Delta'/x\right], \overline{x_{:\omega}\tau}, \overline{y_{i:\Delta_i\left[\Delta'/x\right]}\sigma_i}; \Delta, \Delta' \vdash e\left[e'/x\right] : \varphi$$
 by i.h.(1)

(6)
$$\overline{\Delta_i[\Delta'/x]} = \overline{\Delta_s[\Delta'/x] \# K_i}^n$$
 by (3) and cong.

(7)
$$\Gamma\left[\Delta'/x\right]$$
; $\Delta, \Delta' \vdash \rho \Rightarrow e\left[e'/x\right]$; $\frac{z}{\Delta_{s}\left[\Delta'/x\right]} \varphi\left[\ddagger\right]$ by (5,6)

Case: *Alt*0 This is one of the most interesting proof cases, and challenging to prove.

- The first insight is to divide the proof into two subcases, accounting for when the scrutinee (and hence Δ_s) contains the linear resource and when it does not.
- The second insight is to recall that Δ and Δ' are disjoint to be able to prove the subcase in which x does not occur in the scrutinee
- The third insight is to *create* linear resources seemingly out of nowhere *under a substitution* that removes them. We see this happen in the case where x occurs in the scrutinee, for both the linear and affine contexts (see (5,6)). We must also see that we can swap x for Δ' if neither can occur (see (7)).

(1)
$$\Gamma$$
; $\Delta' \vdash e' : \sigma$

Subcase *x* occurs in scrutinee

Subcase
$$x$$
 occurs in scrutinee

(2) $\Gamma; \Delta, x:_1\tau \vdash K \overline{x:_\omega\tau} \Rightarrow e :_{\Delta_s, x:_1\tau}^z \varphi$

(2.5) $\Gamma \left[\cdot / \Delta_s, x \right]_z, \overline{x:_\omega\tau}; (\Delta, x:_1\tau) \left[\cdot / \Delta_s, x \right] \vdash e : \varphi$ by inv.

(3) $\Gamma \left[\cdot / \Delta_s, x \right]_z, \overline{x:_\omega\tau}; \Delta \left[\cdot / \Delta_s \right] \vdash e : \varphi$

(4) $e[e'/x] = e$ since x cannot occur in e (erased from e cx)

(5) $\Delta \left[\cdot / \Delta_s \right] = (\Delta, \Delta') \left[\cdot / \Delta_s, \Delta' \right]$ by cong. of subst.

(6) $\Gamma \left[\cdot / \Delta_s, x \right]_z \left[\Delta'/x \right] = \Gamma \left[\Delta'/x \right] \left[\cdot / \Delta_s, \Delta' \right]_z$ by cong. of subst.

(7) $\forall x, \Delta, \Delta', \Gamma : x \notin \Delta \land \Delta' \not\subset \Delta \land \Gamma; \Delta \vdash e : \sigma \Rightarrow \Gamma \left[\Delta'/x \right]; \Delta \vdash e : \sigma$ by Weaken and variables in Γ cannot occur in e if they mention x nor if they mention Δ'

(8) $\Gamma \left[\Delta'/x \right] \left[\cdot / \Delta_s, \Delta' \right]_z, \overline{x:_\omega\tau}; (\Delta, \Delta') \left[\cdot / \Delta_s, \Delta' \right] \vdash e[e'x] : \varphi$ by $(4,5,6,7)$ and x and Δ' are erased from ctx

(9) $\Gamma \left[\Delta'/x \right]; \Delta, \Delta' \vdash_{alt} K \overline{x:_\omega\tau} \Rightarrow e[e'/x] :_{\Delta_s,\Delta'}^z \varphi$ by $Alt0$ Subcase x does not occur in scrutinee

(2) $\Gamma; \Delta, x:_1\tau \vdash K \overline{x:_\omega\tau} \Rightarrow e :_{\Delta_s}^z \varphi$

(3) $\Gamma \left[\cdot / \Delta_s \right]_z, \overline{x:_\omega\tau}; \Delta \left[\cdot / \Delta_s \right], x:_1\tau \vdash e : \varphi$ by x does not occur in Δ_s and inv.

(4) $\Gamma \left[\Delta'/x \right] \left[\cdot / \Delta_s \right]_z, \overline{x:_\omega\tau}; \Delta \left[\cdot / \Delta_s \right], \Delta' \vdash e[e'/x] : \varphi$ by i.h.(1) and x does not occur in Δ_s

(5) $\Gamma \left[\Delta'/x \right] \left[\cdot / \Delta_s \right]_z, \overline{x:_\omega\tau}; \Delta \left[\cdot / \Delta_s \right], \Delta' \vdash e[e'/x] : \varphi$

by Δ and Δ' being disjoint by hypothesis, and Δ_s being a subset of Δ

(6)
$$\Delta_s[\Delta'/x] = \Delta_s$$
 by x does not occur in Δ_s

(7) $\Gamma \left[\Delta'/x \right] ; \Delta, \Delta' \vdash K \overline{x :_{\omega} \tau} \implies e[e'/x] :_{\Delta \left[\Delta'/x \right]}^{z} \varphi$

The substitution lemma for unrestricted variables follows the usual formulation, with the added restriction (common to linear type systems) that the expression e' that is going to substitute the unrestricted variable *x* is typed on an empty linear environment:

LEMMA B.5 (SUBSTITUTION OF UNRESTRICTED VARIABLES PRESERVES TYPING).

(1) If
$$\Gamma, x:_{\omega}\sigma; \Delta \vdash e : \varphi \text{ and } \Gamma; \cdot \vdash e' : \sigma \text{ then } \Gamma, \Delta \vdash e[e'/x] : \varphi$$

(2) If $\Gamma, x:_{\omega}\tau; \Delta \vdash \rho \Rightarrow e :_{\Delta}^{z} \sigma \text{ and } \Gamma; \Delta \vdash e' : \tau \text{ then } \Gamma; \Delta \vdash \rho \Rightarrow e[e'/x] :_{\Delta}^{z} \sigma$

PROOF. By structural induction on the given derivation. We omit many trivial cases. Statement (1):

Case: λI_1

$$(1) \Gamma, x:_{\omega}\tau; \Delta \vdash \lambda y:_{1}\sigma'. \ e : \sigma' \rightarrow_{1} \varphi$$

$$(2) \Gamma; \cdot \vdash e' : \sigma$$

$$(3) \Gamma, x:_{\omega}\tau; \Delta, y:_{1}\sigma' \vdash e : \varphi$$

$$(4) \Gamma; \Delta, y:_{1}\sigma' \vdash e[e'/x] : \varphi$$

$$(5) \Gamma; \Delta \vdash \lambda y:_{1}\sigma'. \ e[e'/x] : \sigma' \rightarrow_{1} \varphi$$

$$(6) (\lambda y:_{\pi}\sigma'. \ e)[e'/x] = (\lambda y:_{\pi}\sigma'. \ e[e'/x])$$
by inversion on λI_{1}
by i.h.(1) (2,3)
by λI_{1}
by def. of subst.

Case: Var_{ω}

(1)
$$\Gamma, x:_{\omega}; \cdot \sigma \vdash x : \sigma$$

(2) $\Gamma; \cdot \vdash e' : \sigma$
(4) $x[e'/x] = e'$

by def. of substitution

(5)
$$\Gamma$$
; $\vdash e' : \sigma$ by (2)

Case: Var_{ω}

- (1) Γ , $x:_{\omega}\tau$; $\cdot \vdash y: \varphi$
- (2) Γ : $\cdot \vdash e' : \sigma$
- (3) y[e'/x] = yby def. of substitution
- (4) Γ ; $\cdot \vdash y : \varphi$ by inversion on $Weaken_{\omega}$ (1)

Case: Var₁

(1) Impossible. The context in Var_1 is empty.

Case: Var_{Λ}

(1) Impossible. The context in Var_{Λ} only contains linear variables.

Case: $\lambda E_{1,p}$

- (1) Γ , x: $_{\omega}\tau$; Δ , $\Delta' \vdash e e'' : \varphi$
- (2) Γ : $\cdot \vdash e' : \sigma$
- (3) $\Gamma, x:_{\omega}\tau; \Delta \vdash e : \sigma' \rightarrow_{1,p} \varphi$ by inversion on $\lambda E_{1,p}$ (4) Γ , $x:_{\omega}\tau$; $\Delta' \vdash e'' : \sigma'$ by inversion on $\lambda E_{1,p}$ by i.h.(1) (2,3)
- (5) Γ ; $\Delta \vdash e[e'/x] : \sigma' \rightarrow_{1,p} \varphi$

(6) Γ ; $\Delta' \vdash e''[e'/x] : \sigma'$

by i.h.(1) (2,4)

(7) Γ ; Δ , $\Delta' \vdash e[e'/x] e''[e'/x] : \varphi$

by $\lambda E_{1,p}$ (5,6)

(8) $(e \ e'')[e'/x] = (e[e'/x] \ e''[e'/x])$

by def. of subst.

Case: Let

- (1) Γ , $x:_{\omega}\tau$; Δ , $\Delta' \vdash \text{let } y:_{\Delta}\sigma' = e \text{ in } e'' : \varphi$
- (2) Γ ; $\cdot \vdash e' : \sigma$
- (3) Γ , $x:_{\omega}\tau$, $y:_{\Lambda}\sigma'$; Δ , $\Delta' \vdash e''\varphi$ by inversion on Let by inversion on Let
- (4) $\Gamma, x:_{\omega} \tau; \Delta \vdash e : \sigma'$
- (5) Γ , $y:_{\Delta}\sigma'$; $\Delta \vdash e''[e'/x] : \varphi$ by i.h.(1) (2,3)
- (6) Γ ; $\Delta \vdash e[e'/x] : \sigma'$

by i.h.(1) (2,4)

(7) Γ ; Δ , $\Delta' \vdash \text{let } y: \wedge \sigma' = e[e'/x] \text{ in } e''[e'/x]$

- by Let (5,6)
- (8) (let $y:_{\Delta}\sigma' = e$ in e'')[e'/x] = (let $y:_{\Delta}\sigma' = e[e'/x]$ in e''[e'/x])

Case: LetRec

- (1) $\Gamma, x:_{\omega} \tau; \Delta, \Delta' \vdash \text{let rec } \overline{y:_{\Delta} \sigma' = e} \text{ in } e'' : \varphi$
- (2) Γ' : $\cdot \vdash e' : \sigma$
- (3) $\Gamma, x:_{\omega} \tau, \overline{y:_{\Delta} \sigma'}; \Delta, \Delta' \vdash e'' : \varphi$

by inversion on *LetRec*

(4) $\Gamma, x:_{\omega}\tau, \overline{y:_{\Lambda}\sigma'}; \Delta \vdash e : \sigma'$ (5) Γ , $\overline{y:_{\Delta}\sigma'}$; Δ , $\Delta' \vdash e''[e'/x] : \varphi$

by inversion on *LetRec*

(6) Γ , $\overline{y:_{\Delta}\sigma'}$; $\Delta \vdash e[e'/x] : \sigma'$

by i.h.(1) (2,3) by i.h.(1) (2,4)

(7) Γ ; Δ , $\Delta' \vdash \text{let rec } \overline{y_{:\Delta}\sigma' = e[e'/x]} \text{ in } e''[e'/x] : \varphi$

- by LetRec (5,6)
- (8) (let rec $\overline{y:\Delta\sigma'=e}$ in e'')[e'/x] = (let rec $\overline{y:\Delta\sigma'=e[e'/x]}$ in e''[e'/x])

Statement (2):

Case: $Alt N_{WHNF}$ (trivial induction)

- (1) Γ ; $\cdot \vdash e : \sigma$
- (2) $\Gamma, x:_{\omega} \tau; \Delta \vdash_{alt} K \overline{x:_{\omega} \tau}, \overline{y_{i}:_{1} \sigma_{i}}^{n} \rightarrow e^{\dagger} :_{\underline{A}^{-n}} \sigma' \Longrightarrow \varphi$

$$(3) \Gamma, x:_{\omega} \tau, \overline{x}:_{\omega} \overline{\tau}, \overline{y}_{i}:_{\Delta_{i}} \overline{\sigma}_{i}; \Delta \vdash e : \varphi$$

$$(4) \Gamma, \overline{x}:_{\omega} \overline{\tau}, \overline{y}_{i}:_{\Delta_{i}} \overline{\sigma}_{i}; \Delta \vdash e [e'/x] : \varphi$$

$$(5) \Gamma; \Delta \vdash_{alt} K \overline{x}:_{\omega} \overline{\tau}, \overline{y}_{i}:_{\sigma} \overline{\tau}^{n} \rightarrow e [e'/x]^{\dagger}:_{\overline{\Delta}_{i}} \sigma' \Longrightarrow \varphi$$
by i.h.(1)
by $Alt N$

Case: AltN_{NotWHNF} (trivial induction)

```
 \begin{array}{ll} (1) \; \Gamma; \; \vdash e : \sigma \\ (2) \; \Gamma, x :_{\omega} \tau; \Delta \vdash_{alt} K \; \overline{x :_{\omega} \tau}, \overline{y_{i} :_{1} \sigma_{i}}^{n} \; \to e^{\frac{t}{z}} :_{\Delta_{s}}^{z} \sigma' \; \Rrightarrow \; \varphi \\ (3) \; \overline{\Delta_{i}} \; = \; \overline{\Delta_{s} \# K_{j}}^{n} \\ (4) \; \Gamma, x :_{\omega} \tau, \overline{x :_{\omega} \tau}, \overline{y_{i} :_{\Delta_{i}} \sigma_{i}}; \Delta \vdash e : \varphi \\ (5) \; \Gamma, \overline{x :_{\omega} \tau}, \overline{y_{i} :_{\Delta_{i}} \sigma_{i}}; \Delta \vdash e [e'/x] : \varphi \\ (6) \; \Gamma; \Delta \vdash_{alt} K \; \overline{x :_{\omega} \tau}, \overline{y_{i} :_{1} \sigma_{i}}^{n} \; \to e [e'/x]^{\frac{t}{z} :_{\Delta}^{z}} \; \sigma' \; \Rrightarrow \; \varphi \\ \end{array} \qquad \qquad \qquad \text{by i.h.(1)}
```

Finally, we introduce the lemma stating that substitution of Δ -bound variables by expressions of the same type preserves the type of the original expression. What distinguishes this lemma from traditional substitution lemmas is that the usage environment Δ of the variable x being substituted by expression e' must match exactly the typing environment Δ of e' and the environment of the original expression doesn't change with the substitution:

Lemma B.6 (Substitution of Δ -variables preserves typing).

```
(1) If \Gamma, x:_{\Delta}\sigma; \Delta, \Delta' \vdash e : \varphi \text{ and } \Gamma; \Delta \vdash e' : \sigma \text{ then } \Gamma; \Delta, \Delta' \vdash e[e'/x] : \varphi

(2) If \Gamma, x:_{\Delta}\sigma; \Delta, \Delta' \vdash \rho \Rightarrow e :_{\Delta_s}^{z} \sigma' \text{ and } \Gamma; \Delta \vdash e' : \sigma \text{ and } \Delta_s \subseteq (\Delta, \Delta') \text{ then } \Gamma; \Delta, \Delta' \vdash \rho \Rightarrow e[e'/x] :_{\Delta_s}^{z} \sigma'
```

Intuitively, if x is well-typed with Δ in e, substituting x by an expression e' which is typed in the same environment Δ allows the distribution of resources Δ , Δ' used to type e across sub-derivations to remain unchanged. To prove the theorems, we don't need a "stronger" substitution of Δ -vars lemma (allowing arbitrary resources Δ'' to type e', as in other substitution lemmas), as we only ever substitute Δ -variables by expressions whose typing environment matches the variables usage environment. However, it is not obvious whether such a lemma is possible to prove for Δ -variables (e.g. let Γ ; $\Delta \vdash e : \sigma$ and Γ ; $\Delta' \vdash \operatorname{let} x = e'$ in x, if we substitute e for x the resources Δ' are no longer consumed).

PROOF. By structural induction on the given derivation. We omit many trivial cases. Statement (1):

Case: ΛI

```
(1) \Gamma, x:_{\Delta}\sigma; \Delta, \Delta' \vdash \Lambda p. \ e : \forall p. \ \varphi
(2) \Gamma; \Delta \vdash e' : \sigma
(3) \Gamma, p, x:_{\Delta}\sigma; \Delta, \Delta' \vdash e : \varphi
(3) \Gamma, p, x:_{\Delta}\sigma; \Delta, \Delta' \vdash e : \varphi
(4) \Gamma, p; \Delta, \Delta' \vdash e[e'/x]
(5) \Gamma; \Delta, \Delta' \vdash \Lambda p. \ e[e'/x] : \forall p. \ \varphi
(6) (\Lambda p. \ e)[e'/x] = (\Lambda p. \ e[e'/x])
(7) \Gamma; \Delta, \Delta' \vdash (\Lambda p. \ e)[e'/x] : \forall p. \ \varphi
(8) \text{by def. of subst.}
(9) \Gamma; \Delta, \Delta' \vdash (\Lambda p. \ e)[e'/x] : \forall p. \ \varphi
(9) (5,6)
```

Case: λI_1

zy Linearity for a Core Functional Language	69:41
(1) $\Gamma, x:_{\Delta} \sigma; \Delta, \Delta' \vdash \lambda y:_{1} \sigma'. e : \sigma' \rightarrow_{1} \varphi$	
(2) Γ ; $\Delta \vdash e' : \sigma$	
(3) Γ , $x:_{\Delta}\sigma$; Δ , $y:_{1}\sigma'$, $\Delta' \vdash e : \varphi$	by inversion on λI
(4) Γ ; Δ , y : $_{1}\sigma'$, $\Delta' \vdash e[e'/x] : \varphi$	by i.h.(1) (2,3)
(5) Γ ; Δ , $\Delta' \vdash \lambda y : 1\sigma'$. $e[e'/x] : \sigma' \rightarrow 1\varphi$	by λI
(6) $(\lambda y:_1\sigma'. e[e'/x]) = (\lambda y:_1\sigma'. e)[e'/x]$	by def. of subst.
$(7) \Gamma; \Delta, \Delta' \vdash \lambda(y; \sigma', e)[e'/x] : \sigma' \to_1 \varphi$	by (4,5)
Case: Var_{ω}	
(1) $\Gamma, y:_{\omega}\sigma', x::\sigma; \cdot \vdash y:\sigma'$	
(2) Γ ; \vdash e' : σ	
(3) y[e'/x] = y	
$(4) \Gamma, y:_{\omega} \sigma'; \cdot \vdash y: \sigma'$	by (1) and $Weaken_{\Delta}$
Case: Var_1	
$(1) \Gamma, x_{:y:\tau}\tau; y_{:\tau}\tau \vdash y : \sigma$	
(2) Γ ; $y:_1\tau \vdash e':\sigma$	
$(3) \ y[e'/x] = y$	L., 1
$(4) \Gamma, x:_{y:\tau}\tau; y:_{\tau}\tau \vdash y:\sigma$	by 1 by <i>Weaken</i> ∆
$(5) \Gamma; y:_1 \tau \vdash y : \sigma$	by $weaken_{\Delta}$
Case: Var_{Δ}	
(1) Γ , $x:\Delta \sigma$; $\Delta \vdash y:\sigma$	
(2) Γ' ; $\Delta \vdash e' : \sigma$	
(3) x[e'/x] = e'	
(4) Γ' ; $\Delta \vdash e' : \sigma$	by (2)
Const. 1E	
Case: $\lambda E_{1,p}$	
$(1) \Gamma, x:_{\Delta} \sigma; \Delta, \Delta', \Delta'' \vdash e \ e'' : \varphi$	
(2) Γ ; $\Delta \vdash e' : \sigma$ Subcase Δ occurs in e	
(3) $\Gamma, x:_{\Delta}\sigma; \Delta, \Delta' \vdash e : \sigma' \rightarrow_{1,p} \varphi$	
$(3) \Gamma, \chi_{\cdot \Delta} \sigma, \Delta, \Delta \vdash e \cdot \sigma \longrightarrow_{1,p} \psi$ $(4) \Gamma, \chi_{\cdot \Delta} \sigma, \Delta'' \vdash e'' : \sigma'$	
$(5) \Gamma; \Delta'' \vdash e'' : \sigma'$	by Weaken∧
(6) Γ ; Δ , $\Delta' \vdash e[e'/x] : \sigma' \rightarrow_{1,p} \varphi$	by i.h.(1)
(7) Γ ; Δ , Δ' , $\Delta'' \vdash e[e'/x] e'' : \varphi$	by $(\lambda E_{1,p})$
(8) $(e[e'/x] e'') = (e e'')[e'/x]$	since x cannot occur in e''
Subcase \triangle occurs in e''	
(3) $\Gamma, x:_{\Delta}\sigma; \Delta' \vdash e : \sigma' \rightarrow_{1,p} \varphi$	
$(4) \Gamma; \Delta' \vdash e : \sigma' \rightarrow_{1,p} \varphi$	by Weaken∆
$(5) \Gamma, x:_{\Delta} \sigma; \Delta, \Delta'' \vdash e'' : \sigma'$	
(6) Γ ; Δ , $\Delta'' \vdash e''[e'/x] : \sigma'$	by i.h.(1)
$(7) \Gamma; \Delta, \Delta', \Delta'' \vdash (e \ e''[e'/x]) : \varphi$	by $(\lambda E_{1,p})$
(8) $e e''[e'/x] = (e e'')[e'/x]$	since <i>x</i> does not occur in <i>e</i>
Subcase Δ is split between e and e''	

x cannot occur in either, so the substitution is trivial, and x can be weakened.

Case: λE_{ω}

- (1) Γ , $x: \Lambda \sigma$; Δ , $\Delta' \vdash e e'' : \varphi$
- (2) Γ ; $\Delta \vdash e' : \sigma$
- (3) Δ cannot occur in e''
- (4) Γ , $x:_{\Lambda}\sigma$; Δ , $\Delta' \vdash e : \sigma' \rightarrow_{\omega} \varphi$ by inv. on λE_{ω} (5) Γ ; $\cdot \vdash e'' : \sigma'$ by inv. on λE_{ω} (6) Γ ; Δ , $\Delta' \vdash e[e'/x] : \sigma' \rightarrow_{\omega} \varphi$ by i.h.(1) (2,4)
- (7) Γ ; Δ , $\Delta' \vdash e[e'/x] e'' : \varphi$ by λE_{ω} (5,6)
- (8) e[e'/x] e'' = (e e'')[e'/x]x does not occur in e'' by (3)

Case: Let

(1) Γ ; $\Delta \vdash e' : \sigma$

Subcase Δ occurs in e

- (2) Γ , $x:_{\Delta}\sigma$; Δ , Δ' , $\Delta'' \vdash \text{let } y:_{\Delta,\Delta'}\sigma' = e \text{ in } e'' : \varphi$
- (3) Γ , $x:_{\Lambda}\sigma$; Δ , $\Delta' \vdash e : \sigma'$ by inversion on (let)
- $(4) \; \Gamma, x:_{\Delta} \sigma, y:_{\Delta, \Delta'} \sigma'; \Delta, \Delta', \Delta'' \vdash e'' : \varphi$ by inversion on (let)
- (5) Γ , $y:_{\Delta,\Delta'}\sigma'$; Δ , Δ' , $\Delta'' \vdash e'' : \varphi$ by $Weaken_{\Lambda}$ (admissible)
- (6) Γ ; Δ , $\Delta' \vdash e[e'/x] : \sigma'$ by i.h.(1) (1,3) by (let) (5,6)
- (7) Γ ; Δ , Δ' , $\Delta'' \vdash \text{let } y:_{\Delta,\Delta'} \sigma' = e[e'/x] \text{ in } e'' : \varphi$
- (8) let $y:_{\Delta,\Delta'}\sigma' = e[e'/x]$ in $e'' = (\text{let } y:_{\Delta,\Delta'}\sigma' = e \text{ in } e'')[e'/x]$

since x cannot occur in e''

Subcase Δ occurs in e''

- (2) Γ , $x:\Delta\sigma$; Δ , Δ' , $\Delta'' \vdash \text{let } y:\Delta'\sigma' = e \text{ in } e'' : \varphi$
- (3) Γ , $x: \Lambda \sigma$; $\Delta' \vdash e : \sigma'$ by inversion on (let) (4) Γ ; $\Delta' \vdash e : \sigma'$ by $Weaken_{\Delta}$
- (5) Γ , $x:_{\Delta}\sigma$, $y:_{\Delta'}\sigma'$; Δ , Δ' , $\Delta'' \vdash e'' : \varphi$ by inversion on (let)
- (6) Γ , $y:_{\Delta'}\sigma'$; Δ , Δ' , $\Delta'' \vdash e''[e'/x] : \varphi$ (7) Γ ; Δ , Δ' , $\Delta'' \vdash$ **let** y: $_{\Delta'}\sigma' = e$ **in** $e''[e'/x] : \varphi$
- by i.h.(1) (1,5) by (let)
- (8) let $y:_{\Delta'}\sigma' = e$ in $e''[e'/x] = (\text{let } y:_{\Delta'}\sigma' = e$ in e'')[e'/x]

since x cannot occur in e

Subcase Δ is split between e and e''

x cannot occur in either, so the substitution is trivial, and x can be weakened.

Case: LetRec

(1) Γ ; $\Delta \vdash e' : \sigma$

Subcase Δ occurs in $\overline{e_i}$

- (2) $\Gamma, x:_{\Delta} \sigma; \Delta, \Delta', \Delta'' \vdash \mathbf{let} \ \mathbf{rec} \ \overline{y_{i:_{\Delta,\Delta'}} \sigma'_i = e_i} \ \mathbf{in} \ e'' : \varphi$
- (3) Γ , $x:_{\Delta}\sigma$, $\overline{y_i:_{\Delta,\Delta'}\sigma'_i}$; Δ , Δ' , $\Delta'' \vdash e'' : \varphi$ by inversion on (let)
- (4) Γ , $\overline{y_i:\Delta \Delta' \sigma'_i}$; Δ , Δ' , $\Delta'' \vdash e'' : \varphi$ by Weaken∧
- (5) Γ , $x:_{\Delta}\sigma$, $\overline{y_i:_{\Delta,\Delta'}\sigma'_i}$; Δ , $\Delta' \vdash e_i : \sigma'_i$ by inversion on (let)
- (6) Γ , $\overline{y_i:_{\Delta,\Delta'}\sigma_i'}$; Δ , $\Delta' \vdash e_i[e'/x] : \sigma_i'$ by i.h.(1) (1,5) (7) e''[e'/x] = e''since x cannot occur in e''
- (8) Γ ; Δ , Δ' , $\Delta'' \vdash$ let rec $\overline{y_{i:\Delta,\Delta'}\sigma_i' = e_i[e'/x]}$ in $e'' : \varphi$ by (let) (4,6)

Subcase Δ occurs in e''

- (2) $\Gamma, x:_{\Delta}\sigma; \Delta, \Delta', \Delta'' \vdash \text{let rec } \overline{y_{i:_{\Delta'}}\sigma'_{i} = e_{i}} \text{ in } e'': \varphi$
- (3) $\Gamma, x:_{\Delta} \sigma, \overline{y_i:_{\Delta'} \sigma_i'}; \Delta' \vdash e_i : \sigma_i'$ by inversion on (let)

$$(4) \overline{\Gamma, y_{i:\Delta'}\sigma'_{i}}\Delta' \vdash e_{i}:\sigma'_{i} \qquad \qquad \text{by } Weaken_{\Delta} \\ (6) \Gamma, x_{i\Delta}\sigma_{i}; \Delta, \Delta', \Delta'' \vdash e'':\varphi \qquad \qquad \qquad \text{by inversion on (let)} \\ (7) \Gamma, y_{i:\Delta'}\sigma'_{i}; \Delta, \Delta', \Delta'' \vdash e'' \mid e'/x \mid :\varphi \qquad \qquad \qquad \text{by in.} (1) (1,6) \\ (8) \overline{e_{i}[e'/x]} = e_{i} \qquad \qquad \qquad \text{by in.} (1) (1,6) \\ (9) \Gamma; \Delta, \Delta', \Delta'' \vdash \text{let rec } y_{i:\Delta'}\sigma'_{i} = e_{i} \text{ in } e'' \mid e'/x \mid :\varphi \qquad \qquad \text{by (let)} \\ \text{Subcase } \Delta \text{ is split between } e \text{ and } e'' \qquad \qquad x \text{ cannot occur in } either, so \text{ the substitution is trivial, and } x \text{ can be weakened.} \\ \textbf{Case: CaseWHNF} \\ (1) \Gamma; \Delta \vdash e' : \sigma \qquad \qquad \text{Subcase } \Delta \text{ occurs in } e \\ (2) \Gamma, x_{i:\Delta}\sigma; \Delta, \Delta', \Delta'' \vdash \text{case } e \text{ of } z_{:\Delta,\Delta'}\sigma' \text{ } f \Rightarrow e'' \text{ } } \} \cdot \varphi \\ \text{3) } e \text{ is in WHNF} \\ (4) \rho_{i} \text{ matches } e \\ (5) \Gamma, x_{i:\Delta}\sigma, \Delta, \Delta' \vdash e : \sigma' \qquad \qquad \text{by inv.} \\ (6) \Gamma, x_{i:\Delta}\sigma, z_{:\Delta,\Delta'}\sigma'; \Delta, \Delta', \Delta'' \vdash \rho_{j} \Rightarrow e'' : z_{\Delta,\Delta'}^{2} \varphi \qquad \qquad \text{by inv.} \\ (8) \Gamma; \Delta, \Delta', e \mid e'/x \mid : \sigma' \qquad \qquad \text{by inv.} \\ (8) \Gamma; \Delta, \Delta', e \mid e'/x \mid : \sigma' \qquad \qquad \text{by inv.} \\ (9) \Gamma, x_{:\Delta}\sigma, z_{:\Delta,\Delta'}\sigma'; \Delta, \Delta', \Delta'' \vdash \rho_{j} \Rightarrow e'' : e'/x \mid z_{\Delta,\Delta'}^{2} \varphi \qquad \qquad \text{by } (7) \text{ and proof steps be theneous } from \text{ case } CaseNotWHNF} \\ (11) \Gamma; \Delta, \Delta', \Delta'' \vdash \text{case } e \mid e'/x \mid \sigma' \mid z_{\Delta,\Delta'} \sigma' \mid \overline{\rho} \Rightarrow e'' \mid e'/x \mid z_{\Delta'}^{2} \varphi \qquad \qquad \text{by } (2amissible) Weaken_{\Delta} \\ (12) \Gamma; \Delta, \Delta', \Delta'' \vdash e : \sigma' \qquad \qquad \text{by inv.} \\ (9) \Gamma, x_{:\Delta}\sigma, z_{:\Delta'}\sigma'; \Delta, \Delta', \Delta'' \vdash \rho_{j} \Rightarrow e'' : z_{\Delta}^{2}, \varphi \qquad \qquad \text{by } (admissible) Weaken_{\Delta} \\ (9) \Gamma, x_{:\Delta}\sigma, z_{:\Delta'}\sigma'; \Delta, \Delta', \Delta'' \vdash \rho_{j} \Rightarrow e'' : z_{\Delta}^{2}, \varphi \qquad \qquad \text{by } (admissible) Weaken_{\Delta} \\ (9) \Gamma, x_{:\Delta}\sigma, z_{:\Delta'}\sigma'; \Delta, \Delta', \Delta'' \vdash \rho_{j} \Rightarrow e'' : z_{\Delta}^{2}, \varphi \qquad \qquad \text{by inv.} \\ (10) \Gamma, z_{:\Delta'}\sigma'; \Delta, \Delta', \Delta'' \vdash \rho_{j} \Rightarrow e'' : z_{\Delta}^{2}, \varphi \qquad \qquad \text{by } (admissible) Weaken_{\Delta} \\ (9) \Gamma, x_{:\Delta}\sigma, z_{:\Delta'}\sigma'; \Delta, \Delta', \Delta'' \vdash \rho_{j} \Rightarrow e'' : z_{\Delta}^{2}, \varphi \qquad \qquad \text{by inv.} \\ (10) \Gamma, z_{:\Delta'}\sigma'; \Delta, \Delta', \Delta'' \vdash \rho_{j} \Rightarrow e'' : z_{\Delta}^{2}, \varphi \qquad \qquad \text{by inv.} \\ (10) \Gamma, z_{:\Delta'}\sigma'; \Delta, \Delta', \Delta'' \vdash \rho_{j} \Rightarrow e'' : e'/x \mid z_{\Delta'}^{2}, \varphi \qquad \qquad \text{by inv.} \\ (10) \Gamma, z_{:\Delta'}\sigma'; \Delta, \Delta', \Delta'' \vdash \rho_{j} \Rightarrow e'' : e'/x \mid z_{\Delta'}^{2}, \varphi \qquad \qquad \text{by inv.} \\ (10) \Gamma, z_{:\Delta'}\sigma'; \Delta, \Delta',$$

Case: CaseNotWHNF

(1)
$$\Gamma$$
; $\Delta \vdash e' : \sigma$
Subcase Δ occurs in e

(2)
$$\Gamma, x:_{\Delta}\sigma; \Delta, \Delta', \Delta'' \vdash \text{case } e \text{ of } z:_{[\Delta,\Delta']}\sigma' \{\overline{\rho \Rightarrow e''}\}$$

to contain some of part of Δ and Δ to refer to the other part only.

This is like the subcase above, but consider Δ'

(3)
$$\Gamma, x:_{\Delta} \sigma; \Delta, \Delta' \vdash e : \sigma'$$
 by inv.
(4) $\Gamma; \Delta, \Delta' \vdash e[e'/x] : \sigma'$ by i.h.(1)

(5)
$$\Gamma, x_{:\Delta}\sigma, z_{:[\Delta,\Delta']}\sigma'; [\Delta,\Delta'], \Delta'' \vdash \rho \Rightarrow e'' :_{[\Delta,\Delta']}^z \varphi$$
 by inv.

(6)
$$e''[e'/x] = e$$
 by x cannot occur in $\overline{e''}$ since Δ is not available (only $[\Delta]$) (7) $\overline{\Gamma, z:_{[\Delta, \Delta']}\sigma'; [\Delta, \Delta'], \Delta'' \vdash \rho \Rightarrow e'':_{[\Delta, \Delta']}^z \varphi}$ by (admissible) $Weaken_{\Delta}$ (8) $\Gamma; \Delta, \Delta', \Delta'' \vdash case \ e[e'/x] \ of \ z:_{[\Delta, \Delta']}\sigma' \ \{\overline{\rho \Rightarrow e''}\}$ by $CaseNotWHNF$ Subcase Δ does not occur in e (2) $\Gamma, x:_{\Delta}\sigma; \Delta, \Delta', \Delta'' \vdash case \ e \ of \ z:_{[\Delta']}\sigma' \ \{\overline{\rho \Rightarrow e''}\}$ by $CaseNotWHNF$ (3) $\Gamma, x:_{\Delta}\sigma; \Delta, \Delta', \Delta'' \vdash e: \sigma'$ by inv. (4) $\Gamma; \Delta' \vdash e: \sigma'$ by weaken (5) $e[e'/x] = e$ by $e[e'/x] = e$ by $e[e'/x] = e$ by $e[e'/x] = e$ by inv. (7) $e[e'/x] = e$ by inv. (8) $e[e'/x] = e$ by $e[e'/x] = e$ by inv. (9) $e[e'/x] = e$ by inv. (10) $e[e'/x] = e$ by inv. (11) $e[e'/x] = e$ by $e[e/x] = e$ by

Statement (2):

Case: Alt0

(1) Γ ; $\Delta \vdash e' : \sigma$

Subcase Δ occurs in scrutinee

(2)
$$\Gamma, x:_{\Delta} \sigma; \Delta, \Delta', \Delta'' \vdash K \overline{x:_{1}\tau} \implies e:_{\Delta, \Delta'}^{z} \varphi$$

$$(3)\; (\Gamma, x:_{\Delta}\sigma)[\cdot/\Delta, \Delta']_z; (\Delta, \Delta', \Delta'')[\cdot/\Delta, \Delta'] \vdash e: \varphi$$

This is like the subcase above, but consider Δ'

to contain some of part of Δ and Δ to refer to the other part only.

(4)
$$\Gamma[\cdot/\Delta, \Delta']_z, x:_{\Delta}\sigma; \Delta'' \vdash e : \varphi$$

by def. of $[]_z$ subst. and [] subst.

(5)
$$\Gamma[\cdot/\Delta, \Delta']_z, x: \Lambda \sigma; \Delta'' \vdash e[e'/x] : \varphi$$

by x cannot occur in e by Δ not available

(6)
$$\Gamma[\cdot/\Delta, \Delta']_z; \Delta'' \vdash e[e'/x] : \varphi$$

by (admissible) $Weaken_{\Lambda}$

$$(7) \cdot = (\Delta, \Delta') [\cdot / \Delta, \Delta']$$

(8)
$$\Gamma[\cdot/\Delta, \Delta']_z$$
; $(\Delta, \Delta', \Delta'')[\cdot/\Delta, \Delta'] \vdash e[e'/x] : \varphi$

by (6,7)

(9)
$$\Gamma$$
; Δ , Δ' , $\Delta'' \vdash e[e'/x] :_{\Lambda}^{z}, \varphi$

Subcase Δ does not occur in the scrutinee

(2)
$$\Gamma, x:_{\Delta} \sigma; \Delta, \Delta', \Delta'' \vdash K \overline{x:_{1}\tau} \implies e:_{\Delta'}^{z} \varphi$$

(3)
$$(\Gamma, x:_{\Delta}\sigma)[\cdot/\Delta']_z; (\Delta, \Delta', \Delta'')[\cdot/\Delta'] \vdash e : \varphi$$

(4)
$$\Gamma[\cdot/\Delta']_z$$
, $x:_{\Delta}\sigma$; Δ , $\Delta'' \vdash e: \varphi$

by def. of $[]_z$ subst. and [] subst.

(5)
$$\Gamma[\cdot/\Delta']_z; \Delta, \Delta'' \vdash e[e'/x] : \varphi$$

by i.h.(1)

$$(6) \cdot = \Delta' [\cdot / \Delta']$$

$$(7) \; \Gamma[\cdot/\Delta']_z; (\Delta,\Delta',\Delta'')[\cdot/\Delta'] \vdash e[e'/x] : \varphi$$

by (5,6)

(8)
$$\Gamma$$
; Δ , Δ' , $\Delta'' \mapsto e[e'/x] :_{\Lambda'}^{z} \varphi$

by Alt0

Subcase Δ is partially used in the scrutinee

This is like the subcase above, but consider Δ'

to contain some of part of Δ and Δ to refer to the other part only.

B.4 Type Preservation

```
Theorem B.7 (Type Preservation). Let \Gamma; \Delta \vdash (\Theta \mid e) \Downarrow (\Theta' \mid v) : \tau, \Sigma or \Gamma; \Delta \vdash (\Theta \mid e) \Downarrow^* (\Theta' \mid v) : \tau, \Sigma. If \Gamma; \Delta \vdash \Theta \mid e : \tau, \Sigma then \Gamma; \Delta \vdash \Theta' \mid v : \tau, \Sigma
```

Proof. By structural induction on the instrumented operational semantics. We omit many trivial cases.

```
Case: \Gamma; \Delta \vdash (\Theta, x:_1\tau = y \mid x) \downarrow (\Theta' \mid v) : \sigma, \Sigma
      (1) \Gamma; \Delta \vdash \Theta, x:_1\tau = y \mid x:\sigma, \Sigma
      (2) \Gamma, \Theta \uparrow^{\omega}; \Delta, \Theta \uparrow^{1}, x:_{1}\tau \vdash (x, y, \Theta \downarrow_{e}^{1}, \Sigma) : \sigma \otimes \sigma \bigotimes (\Theta \downarrow_{e}^{1}) \bigotimes (\Sigma)
                                                                                                                                                                                                                                    by A.3
                                                                                                                                                                            s.t. (\Delta_2, \Delta_3) types (\Theta \downarrow_e^1, \Sigma)
      (3) Let \Gamma'; \Delta_1, \Delta_2, \Delta_3 = \Gamma, \Theta \uparrow^{\omega}; \Delta, \Theta \uparrow^1
      (4) \Gamma'; x:_1\tau \vdash x:\sigma
      (5) \Gamma'; \Delta_1 \vdash y : \sigma
      (6) \Gamma'; \Delta_1, \Delta_2, \Delta_3 \vdash (y, \Theta \downarrow_e^1, \Sigma) : \sigma \bigotimes (\Theta \downarrow_e^1) \bigotimes (\Sigma)
      (7) \Gamma; \Delta \vdash \Theta \mid y : \sigma, \Sigma
                                                                                                                                                                                                   by well-typed state
      (8) \Gamma; \Delta \vdash \Theta' \mid v : \sigma, \Sigma
                                                                                                                                                                                     by induction hypothesis
Case: \Gamma; \Delta \vdash (\Theta, x:_{\Delta'}\sigma = e \mid x) \Downarrow (\Theta', x:_{\Delta_a}\sigma = v \mid v) : \sigma, \Sigma
      (1) \Gamma; \Delta \vdash (\Theta, x :_{\Delta'} \sigma = e \mid x) : \sigma, \Sigma
      (2) \Gamma, \Theta \uparrow^{\omega}; \Delta, \Theta \uparrow^{1} \vdash \text{let } x:_{\Delta'} = e \text{ in } (x, \Theta \downarrow_{\rho}^{1}, \Sigma) : \sigma \bigotimes (\Theta \downarrow_{\sigma}^{1}) \bigotimes (\Sigma)
                                                                                                                                                                                                                                    by A.3
      (3) Let \Delta', \Delta'' = \Delta, \Theta \uparrow^1
      (4) \Gamma, \Theta \uparrow^{\omega}; \Delta' \vdash e : \sigma
                                                                                                                                                                                                            by inv on let (2)
      (5) \Gamma, \Theta \uparrow^{\omega}; \Delta', \Delta'' \vdash (x, \Theta \downarrow_{e}^{1}, \Sigma) : \sigma \bigotimes (\Theta \downarrow_{\sigma}^{1}) \bigotimes (\Sigma)
                                                                                                                                                                                                            by inv on let (2)
      (6) \Gamma, \Theta \uparrow^{\omega}; \Delta' \vdash x : \sigma
```

 $(7) \ \Gamma, \Theta \!\! \uparrow^\omega; \Delta'' \vdash (\Theta \!\! \downarrow_e^1, \Sigma) : \bigotimes (\Theta \!\! \downarrow_\sigma^1) \bigotimes (\Sigma) \\ (8) \ \Gamma, \Theta \!\! \uparrow^\omega; \Delta', \Delta'' \vdash (e, \Theta \!\! \downarrow_e^1, \Sigma) : \tau \bigotimes (\Theta \!\! \downarrow_\sigma^1) \bigotimes (\Sigma) \\ (9) \ \Gamma; \Delta \vdash \Theta \mid e : \tau, \Sigma$ by well-typed state

(9) Γ ; $\Delta \vdash \Theta \mid e : \tau$, Σ by induction hypothesis

(11) Let $\Gamma'; \Delta_a, \Delta_b = \Gamma, \Theta' \uparrow^\omega, \Delta, \Theta' \uparrow^1$

 $(11a) \Gamma'; \Delta_a \vdash v : \sigma (11b) \Gamma'; \Delta_b \vdash (\Theta' \downarrow_e^1, \Sigma) : \bigotimes(\Theta' \downarrow_\sigma^1) \bigotimes(\Sigma)$

(12) Γ' ; $\Delta_a \vdash \text{let } x:_{\Delta_a} \sigma = v \text{ in } v : \sigma$ by let (11a,11a)

(13) Γ ; $\Delta \vdash \Theta' \mid \mathbf{let} \ x :_{\Delta_a} \sigma = v \ \mathbf{in} \ v : \sigma, \Sigma$ by well-typed state (14) Γ ; $\Delta \vdash \Theta', x :_{\Delta_a} \sigma = v \mid v : \sigma, \Sigma$ by env. expansion

Note: It's slightly more elaborate to accommodate recursion. The runtime environment bindings must be tagged with the recursive group such that the environment expansion

can decide whether to unwrap to a letrec group or not. Unlike the case proved here, when the Δ -binding belongs to a let group, the whole group is unwrapped and we can invert the assumption to get an e typed with the xs from the group.

Case: Γ ; $\Delta \vdash (\Theta \mid \text{case } e \text{ of } z_{:\Delta'} \sigma \{K \overline{y_i} \Rightarrow e'\}) \Downarrow (\Theta'' \mid v) : \tau$

```
(1) \Gamma; \Delta \vdash \Theta \mid \mathbf{case} \ e \ \mathbf{of} \ z:_{\Delta_a} \sigma \ \{K \ \overline{y_i} \Rightarrow e'\} : \tau, \Sigma

(2) \Gamma, \Theta \uparrow^{\omega}; \Delta, \Theta \uparrow^1 \vdash (\mathbf{case} \ e \ \mathbf{of} \ z:_{\Delta_a} \sigma \ \{K \ \overline{y_i} \Rightarrow e'\}, \Theta \downarrow_{e}^1, \Sigma) : \tau \otimes \bigotimes (\Theta \downarrow_{\sigma}^1) \otimes \bigotimes (\Sigma) by A.3
```

(3) Let
$$\Gamma' = \Gamma$$
, $\Theta \uparrow^{\omega}$; Δ_a , Δ_b , Δ_2 , $\Delta_3 = \Delta$, $\Theta \uparrow^1$ s.t. (Δ_2, Δ_3) types $(\Theta \downarrow_e^1, \Sigma)$

(4) Γ' ; Δ_a , $\Delta_b \vdash \mathbf{case} \ e \ \mathbf{of} \ z:_{\Delta_a} \sigma\{K \ \overline{y_i} \Longrightarrow e'\} : \tau$

Subcase *e* is in WHNF (thus = $K \overline{x_i}$)

(5) Γ' ; $\Delta_a \Vdash K \overline{x_i} : \sigma > \overline{\Delta_i}$ by subcase

where (Δ_4, Δ_5) types $(\Theta' \downarrow_e^1, \Sigma)$

```
(6) \Gamma', z:_{\Delta_a} \sigma; \Delta_a, \Delta_b \vdash_{alt} K \overline{y_i} \to e' :_{\Delta_i}^z \sigma \Rightarrow \tau
                                                                                                                                                                                                 by inv
Subsubcase \vdash Alt_{WHNF}
(7) \Gamma', z:_{\Delta_a} \sigma, \overline{y_i:_{\omega} \sigma_i}, \overline{y_j:_{\Delta_i} \sigma_j}; \Delta_a, \Delta_b \vdash e' : \tau
                                                                                                                                                                                                 by inv
(8) \Theta' = \Theta
                                                                                                     in this subcase because (\Theta \mid K \overline{x_i}) \downarrow (\Theta \mid K \overline{x_i})
(9) \Gamma'; \overline{\Delta_i} \vdash K \overline{x_i} : \sigma
                                                                                                                                                                                         by inv (5)
(10) \overline{\Gamma'; \cdot \vdash x_i : \sigma_i
                                                                                                                                             by inv (9) for unrestricted xs
(11) \overline{\Gamma'; \Delta_i \vdash x_i : \sigma_i}
                                                                                                                                               by inv (9) for linear fields xs
(12) \Gamma', z:_{\Delta_a} \sigma; \Delta_a, \Delta_b \vdash e' \overline{[x_{i,i}/y_{i,i}] : \tau}
                                                                                                                   by Δ-subst and \omega-subst lemmas (7,10,11)
(13) \Gamma'; \Delta_a, \Delta_b \vdash e' \overline{[x_{i,j}/y_{i,j}]} [K \overline{x_{i,j}}/z] : \tau
                                                                                                                           by Δ-subst lemma and \Delta_i = \Delta_a (9,12)
(14) \Gamma; \Delta \vdash \Theta \mid e' \overline{[x_{i,j}/y_{i,j}]} [K \overline{x_{i,j}}/z] : \tau, \Sigma
                                                                                                                                                                    by well-typed state
(15) \Gamma; \Delta \vdash \Theta'' \mid v : \tau, \Sigma
                                                                                                                                                        by induction hypothesis
Subsubcase \vdash Alt_0
(7) (\Gamma', z:_{\Delta_a} \sigma) [\cdot/\Delta_a]_z, \overline{y_i:_{\omega} \sigma_i}; (\Delta_a, \Delta_b) [\cdot/\Delta_a] \vdash e' : \tau
                                                                                                                                                                              by inv subcase
(8) \Gamma', z:, \overline{y_i:_{\omega}\sigma_i}; \Delta_b \vdash e' : \tau
                                                                                                                                                          by z cannot occur in \Gamma'
(9) Since \Gamma'; \Delta_a \Vdash K \overline{x_i}, in this subcase \Delta_a = \cdot
(10) \Gamma', z:_1\tau:_{\sigma}, \overline{y:_{\omega}\sigma}; \Delta_b \vdash e':_{\tau}
                                                                                                                                                                                                  by (9)
(11) \Gamma', z:_1 \tau:_{\cdot} \sigma; \Delta_b \vdash e' \overline{[x_i/y_i]} : \tau
                                                                                                                                                                            by \omega-subst (10)
(12) \Gamma'; \Delta_b \vdash e' \overline{[x_i/y_i]} [K \overline{x_i}/z] : \tau
                                                                                                                                                                                      by \Delta-subst
(13) \Gamma; \Delta \vdash \Theta \mid e' \overline{[x_i/y_i]} [K \overline{x_i}/z] : \tau, \Sigma
                                                                                                                                                                    by well-typed state
(14) \Gamma; \Delta \vdash \Theta'' \mid v : \tau, \Sigma
                                                                                                                                                        by induction hypothesis
Subcase e is not in WHNF
(5) \Gamma'; \Delta_a \vdash e : \sigma
                                                                                                                                                                              by inv on case
(6) \Gamma', z:_{[\Delta_a]}\sigma; [\Delta_a], \Delta_b \vdash_{alt} K \overline{y_{i,j}} \rightarrow e':_{[\Delta_a]}^z \sigma
                                                                                                                                                                              by inv on case
Subsubcase \vdash Alt_{NWHNF}
                                                                                                                                                     for \Delta_j = \overline{[\Delta_a] \# K_j}, by inv.
(7) \Gamma', z:_{[\Delta_a]} \sigma, \overline{y_i:_{\omega} \sigma_i}, \overline{y_j:_{\Delta_i} \sigma_j}; [\Delta_a], \Delta_b \vdash e' : \tau
(8) \Gamma', z:_{\overline{u_i:1\sigma_i}}\sigma, \overline{u_i:_{\omega}\sigma_i}; \overline{u_i:_{1}\sigma_i}, \Delta_b \vdash e' : \tau
                                                                                                                                                               by \Delta-bound to linear
(9) \Gamma', z:_{\overline{y_i:_1\sigma_i}} \sigma, \overline{y_i:_{\omega}\sigma_i}; \Delta_a, \Delta_b, \overline{y_j:_1\sigma_j} \vdash (e, e') : \sigma \otimes \tau
                                                                                                                                                                                    by pair (5,8)
(10) \Gamma, z: \overline{y_i:_1\sigma_i}\sigma, \overline{y_i:_\omega\sigma_i}; \Delta, \overline{y_j:_1\sigma_j} \vdash \Theta \mid e:\sigma, e':\tau, \Sigma
                                                                                                                                                                   by well-typed state
(11) \Gamma, z: \overline{y_{j:1}\sigma_j}\sigma, \overline{y_i:_{\omega}\sigma_i}; \Delta, \overline{y_{j:1}\sigma_j} \vdash \Theta' \mid K \overline{x_i} : \sigma, e' : \tau, \Sigma
                                                                                                                                                        by induction hypothesis
(12) \Gamma, \Theta' \uparrow^{\omega}, \overline{y_{i}:_{\omega}\sigma_{i}}; \Delta, \overline{y_{j}:_{1}\sigma_{j}}, \Theta' \uparrow^{1} \vdash (K \overline{x_{i,j}}, \Theta' \downarrow_{e}^{1}, e', \Sigma) : \sigma \bigotimes(\Theta' \downarrow_{\sigma}^{1}) \otimes \tau \bigotimes(\Sigma)
(13) Let \Gamma'' = \Gamma, \Theta' \uparrow^{\omega} and \Delta_c \Delta_d \Delta_4 \Delta_5 = \Delta, \Theta' \uparrow^1
                                                                                                                                          where (\Delta_4, \Delta_5) types (\Theta' \downarrow_e^1, \Sigma)
(13a) \Gamma''; \Delta_c \vdash K \overline{x_{i,j}} : \sigma
(13b) \Gamma'', z:_{\overline{y_j:_1\sigma_j}}\sigma, \overline{y_i:_\omega\sigma_i}; \Delta_d, \overline{y_j:_1\sigma_j} \vdash e': \tau
(14) \Gamma''; \cdot \vdash x_i : \sigma_i
                                                                                                                                                                   for \overline{D_i} = \Delta_c, by inv
(15) \Gamma''; \Delta_i \vdash x_i : \sigma_i
(16) \Gamma'', z:_{\Delta_c} \sigma; \Delta_c, \Delta_d \vdash e'[x_{i,j}/y_{i,j}] : \tau
                                                                                                                                    by 1-substs and \omega-substs lemmas
(17) \Gamma''; \Delta_c, \Delta_d \vdash e'[x_{i,j}/y_{i,j}][K\overline{x_{i,j}}/z] : \tau
                                                                                                                                                   by \Delta-subst lemma (13a,16)
(18) \Gamma; \Delta \vdash \Theta' \mid e' \overline{[x_{i,j}/y_{i,j}]} [K \overline{x_{i,j}}/z] : \tau, \Sigma
                                                                                                                                                                    by well-typed state
(19) \Gamma; \Delta \vdash \Theta'' \mid v : \tau, \Sigma
                                                                                                                                                        by induction hypothesis
Subsubcase \vdash Alt_0
(7) \Gamma', z: \sigma, \overline{y_i:_{\omega}\sigma_i}; \Delta_b \vdash e' : \tau
                                                                                                                                                                                                 by inv
                                                                                                                                                                                   by pair (5,7)
(8) \Gamma', z: \sigma, \overline{y_i:_{\omega}\sigma_i}; \Delta_a, \Delta_b \vdash (e, e') : \sigma \otimes \tau
(9) \Gamma, z:.\sigma, \overline{y_i}:\omega \sigma_i; \Delta \vdash \Theta \mid e : \sigma, e' : \tau, \Sigma
                                                                                                                                                                    by well-typed state
(10) \Gamma, z:.\sigma, \overline{y_i}:\omega \sigma_i; \Delta \vdash \Theta' \mid K \overline{x_i} : \sigma, e' : \tau, \Sigma
                                                                                                                                                        by induction hypothesis
(11) Let \Gamma'' = \Gamma, \Theta' \uparrow^{\omega} and \Delta_d, \Delta_4, \Delta_5 = \Delta, \Theta' \uparrow^1
```

(11b) $\Gamma'', z:.\sigma, \overline{y_i:_{\omega}\sigma_i}; \Delta_d \vdash e' : \tau$

(13) $\Gamma'', z: \sigma; \Delta_d \vdash e' \overline{[x_i/y_i]} : \tau$

(11a) Γ'' ; $\cdot \vdash K \overline{x_i} : \sigma$

(12) $\overline{\Gamma''$; $\cdot \vdash x_i : \sigma_i$

by inv (11a)

by ω -subst lemma

since K has no linear fields in this subcase

```
(14) \Gamma''; \Delta_d \vdash e' \overline{[x_i/y_i]} [K \overline{x_i}/z] : \tau
                                                                                                                                                                by Δ-subst lemma
     (15) \Gamma; \Delta \vdash \Theta' \mid e' \overline{[x_i/y_i]} [K \overline{x_i}/z] : \tau, \Sigma
                                                                                                                                                              by well-typed state
     (16) \Gamma; \Delta \vdash \Theta'' \mid v : \tau, \Sigma
                                                                                                                                                   by induction hypothesis
Case: \Gamma; \Delta \vdash (\Theta \mid \mathbf{case} \ e \ \mathbf{of} \ z_{:\Lambda'} \sigma \{\_ \Rightarrow e'\}) \downarrow (\Theta'' \mid v) : \tau
     (1) \Gamma; \Delta \vdash \Theta \mid \mathbf{case} \ e \ \mathbf{of} \ z_{\Delta_a} \sigma \{\_ \Rightarrow e'\} : \tau, \Sigma
     (2) \Gamma, \Theta \uparrow^{\omega}; \Delta, \Theta \uparrow^{1} \vdash (case \ e \ of \ z_{:\Delta_{a}} \sigma \{\_\Rightarrow e'\}, \Theta \downarrow_{e}^{1}, \Sigma) : \tau \otimes \bigotimes(\Theta \downarrow_{\sigma}^{1}) \otimes \bigotimes(\Sigma)
     (3) Let \Gamma' = \Gamma, \Theta \uparrow^{\omega}; \Delta_a, \Delta_b, \Delta_2, \Delta_3 = \Delta, \Theta \uparrow^1
                                                                                                                                           s.t. (\Delta_2, \Delta_3) types (\Theta \downarrow_e^1, \Sigma)
     (4) \Gamma'; \Delta_a, \Delta_b \vdash \mathbf{case} \ e \ \mathbf{of} \ z_{\Delta_a} \sigma \{\_ \Rightarrow e'\} : \tau
     (5) \Gamma'; \Delta_a \vdash e : \sigma
     Subcase e is in WHNF
     (-) Similar to the same subcase in the case-with-pattern-vars reduction,
     except only Alt is relevant which is far simpler
     Subcase e is not in WHNF
     (6) \Gamma', z:_{[\Delta_a]}\sigma; [\Delta_a], \Delta_b \vdash e' : \tau
     (7) \Gamma'; \Delta_b, z:_1 \sigma \vdash e' : \tau
                                                                                                                                                          by \Delta-bound to linear
     (8) \Gamma'; \Delta_a, \Delta_b, z:_1 \sigma \vdash (e, e') : \sigma \otimes \tau
                                                                                                                                                                                        by pair
                                                                                                                                                              by well-typed state
     (9) \Gamma; \Delta, z:<sub>1</sub>\sigma \vdash \Theta \mid e : \sigma, e' : \tau, \Sigma
     (10) \Gamma; \Delta, z:<sub>1</sub>\sigma \vdash \Theta' \mid K \overline{x_i} : \sigma, e' : \tau, \Sigma
                                                                                                                                                   by induction hypothesis
     (11) Let \Gamma'; \Delta_c, \Delta_d, \Delta_4, \Delta_5 = \Gamma, \Theta' \uparrow^{\omega}; \Delta, \Theta' \uparrow^1
                                                                                                                                             s.t. (\Delta_4, \Delta_5) types (\Theta' \downarrow_{\rho}^1, \Sigma)
     (12) \Gamma'; \Delta_c \vdash K \overline{x_i} : \sigma
     (13) \Gamma'; \Delta_d, z:_1 \sigma \vdash e' : \tau
     (14) \Gamma'; \Delta_c, \Delta_d \vdash e'[K \overline{x_i}/z] : \tau
                                                                                                                                                                  by 1-subst lemma
     (15) \Gamma; \Delta \vdash \Theta' \mid e'[K \overline{x_i}/z] : \tau, \Sigma
                                                                                                                                                              by well-typed state
     (16) \Gamma; \Delta \vdash \Theta'' \mid v : \tau, \Sigma
                                                                                                                                                   by induction hypothesis
```

We also derive type preservation for the lazy natural semantics.

Theorem B.8 (Type Preservation). For any well-typed $[\Theta \mid e]$, if $\Theta : e \Downarrow \Theta' : e'$ or $\Theta : e \Downarrow^* \Theta' : e'$ then we have that $[\Theta \mid e']$ is well-typed.

PROOF. By Lemma A.5 and Theorem B.7.

B.5 Progress

Theorem B.9 (Progress). Let Γ ; $\Delta \vdash \Theta \mid e : \tau, \Sigma$. Any partial derivation Γ ; $\Delta \vdash (\Theta \mid e) \Downarrow ?, \Sigma$ can be extended.

PROOF. Essentially all cases are standard. The only potential issue is the (linear) variable rule. It suffices to show that whenever a variable is evaluated it is present in the evaluation environment. This situation is impossible since the state Γ ; $\Delta \vdash \Theta \mid x : \tau, \Sigma$ where x is absent from Θ is not well-typed and so cannot be the head of a partial derivation.

Theorem B.10 (Progress). For any partial derivation of $\Theta: e \downarrow ?$, for $[\Theta|e]$ well-typed, the derivation can be extended.

PROOF. Follows from Lemma A.6 and Theorem B.9.

C Optimisations preserve linearity

THEOREM C.1 (INLINING PRESERVES TYPES).

C.1 Inlining

Inlining substitutes occurrences of a let-bound Δ -variable x with the expression e it is bound to, which determines its usage environment Δ . Intuitively, in the let body e', x can occur once or not at all: if x occurs, then the linear resources Δ used indirectly through x are used via e instead; if x does not occur, then the resources Δ are already used linearly in e' and the substitution is a no-op.

```
If \Gamma; \Delta, \Delta' \vdash \mathbf{let} \ x:_{\Delta} \sigma = e \ \mathbf{in} \ e' : \varphi \ then \ \Gamma; \Delta, \Delta' \vdash \mathbf{let} \ x:_{\Delta} \sigma = e \ \mathbf{in} \ e' [e/x] : \varphi
PROOF.

(1) \Gamma; \Delta, \Delta' \vdash \mathbf{let} \ x:_{\Delta} \sigma = e \ \mathbf{in} \ e' : \varphi
(2) \Gamma, \Delta \vdash e : \sigma by inv. on (let)
(3) \Gamma, x:_{\Delta} \sigma; \Delta, \Delta' \vdash e' : \varphi by inv. on (let)
(4) \Gamma; \Delta, \Delta' \vdash e' [e/x] : \varphi by \Delta-subst. lemma (2,3)
(5) \Gamma, x:_{\Delta} \sigma; \Delta, \Delta' \vdash e' [e/x] : \varphi by (admissible) Weaken_{\Delta}
(6) \Gamma; \Delta, \Delta' \vdash \mathbf{let} \ x:_{\Delta} \sigma = e \ \mathbf{in} \ e' [e/x] : \varphi by (let) (2,5)
```

C.2 β -reduction

 β -reduction evaluates a λ -abstraction application by substituting the λ -bound variable x with the argument e' in the body of the λ -abstraction e. We consider two definitions of β -reductions, one that substitutes all occurrences of a variable by an expression, as in call-by-name, and other which creates a lazy let binding to share the result of computing the argument expression amongst uses of the variable, as in call-by-need.

The first kind of β -reduction on term abstractions can be seen to preserve linearity by case analysis. When the function is linear, the binding x is used exactly once in the body of the lambda, thus can be substituted by an expression typed with linear resources, since the expression is guaranteed to be used exactly once in place of x. The proof is direct by type preservation.

```
Theorem C.2 (\beta-reduction preserves types). If \Gamma; \Delta \vdash (\lambda x:_{\pi} \sigma. e) \ e' : \varphi \ then \ \Gamma; \Delta \vdash e[e'/x] : \varphi Proof. (1) \ \Gamma; \Delta \vdash (\lambda x:_{\pi} \tau. e) \ e' : \varphi (2) (\lambda x:_{\pi} \tau. e) e' \longrightarrow e[e'/x] (3) \Gamma; \Delta \vdash e[e'/x] : \varphi by type preservation theorem (1,2)
```

We assume the β -reduction with sharing (i.e. the one that creates a let binding) is only applicable when the λ -abstraction has an unrestricted function type. Otherwise, the call-by-name β -reduction is always favourable, as we know the resource to be used exactly once and hence sharing would be counterproductive, and result in an unnecessary heap allocation. Consequently, the argument to the function must be unrestricted (hence use no linear resources) for the term to be well-typed, and so it vacuously follows that linearity is preserved by this transformation.

Theorem C.3 (β -reduction with sharing preserves types).

```
If \Gamma; \Delta \vdash (\lambda x :_{\omega} \tau . e) e' : \varphi \text{ then } \Gamma; \Delta \vdash \text{let } x = e' \text{ in } e : \varphi
```

NB: We only apply this transformation when x is unrestricted, otherwise beta-reduction without sharing is the optimization applied.

Proof.

```
\begin{array}{lll} (1) \ \Gamma; \Delta \vdash (\lambda x:_{\omega} \tau. \ e) \ e' : \varphi \\ \\ (2) \ \Gamma; \Delta \vdash (\lambda x:_{\omega} \tau. \ e) : \sigma \rightarrow_{\omega} \varphi & \text{by inv. on } \lambda E_{\omega} \\ \\ (3) \ \Gamma; \cdot \vdash e' : \sigma & \text{by inv. on } \lambda E_{\omega} \\ \\ (4) \ \Gamma, x:_{\omega} \tau; \Delta \vdash e : \varphi & \text{by inv. on } \lambda I \\ \\ (5) \ \Gamma, x:.\sigma; \Delta \vdash e : \varphi & \text{by Lemma 4.3} \\ \\ (6) \ \Gamma, \Gamma' \vdash \mathbf{let} \ x = e' \ \mathbf{in} \ e : \varphi & \text{by let (5,3)} \end{array}
```

Finally, β -reduction on multiplicity abstractions is also type preserving. The argument of the application is a multiplicity rather than an expression, so no resources are needed to type it, and, since the body of the lambda must treat the multiplicity p as though it were linear, the body uses the argument linearly regardless of the instantiation of p at π .

```
Theorem C.4 (\beta-reduction on multiplicity abstractions preserves types). If \Gamma; \Delta \vdash (\Lambda p. e) \pi : \varphi then \Gamma; \Delta \vdash e[\pi/p] : \varphi
```

Proof. Trivial by invoking preservation using the Λ application reduction and the assumption

C.3 Case of known constructor

Case-of-known constructor is a transformation that essentially evaluates a case expression of a known constructor at compile time, by substituting in the matching alternative the case binder by the scrutinee and pattern variables by constructor arguments. Intuitively, either the case binder is used exactly once to consume the resources of the scrutinee, or the pattern components whose matching constructor argument uses linear resources are used exactly once, so one of the substitutions is a no-op. If the pattern variables are substituted by the matching scrutinee expressions, the expressions are still only used once, and if the case binder is substituted by the scrutinee, it is still used exactly once. The proof follows trivially from the preservation theorem.

```
Theorem C.5 (Case-of-known-constructor preserves types). If \Gamma; \Delta, \Delta' \vdash case K \overline{e} of z:_{\Delta}\sigma \{..., K \overline{x} \Rightarrow e_i\} : \varphi then \Gamma; \Delta, \Delta' \vdash e_i [e/x] [K \overline{e}/z] : \varphi Proof.
```

(1)
$$\Gamma; \Delta, \Delta' \vdash \mathbf{case} \ K \ \overline{e} \ \mathbf{of} \ z:_{\Delta} \sigma \ \{..., K \ \overline{x} \Rightarrow e_i\} : \varphi$$
(2) $\mathbf{case} \ K \ \overline{e} \ \mathbf{of} \ z:_{\Delta} \sigma \ \{..., K \ \overline{x} \Rightarrow e_i\} \longrightarrow e_i \overline{[e/x]} [K \ \overline{e}/z]$
(3) $\Gamma; \Delta, \Delta' \vdash e_i \overline{[e/x]} [K \ \overline{e}/z] : \varphi$ by preservation theorem (1,2)

by inv.

C.4 Case of Case

The case of case transformation applies to case expressions whose scrutinee is another case expression, and returns the innermost case expression transformed by repeating the outermost case expression in each alternative of the innermost case, scrutinizing the original alternative body.

Intuitively, since the scrutinee of the outermost case is not in WHNF, no resources from it can directly occur in the outermost alternatives. By moving the outermost alternatives inwards with a different scrutinee, the alternatives remain well-typed because they are typed using either the case binder or the pattern bound variables, which, by the *Irrelevance* lemma, makes it well-typed for any scrutinee consuming arbitrary resources.

Theorem C.6 (Case-of-case preserves types). If Γ ; Δ , Δ' , Δ'' \vdash case case e_c of z: $_{\Delta}\tau$ $\{\overline{\rho_{c_i} \Rightarrow e_{c_i}}\}$ of w: $_{[\Delta,\Delta']}\sigma'$ $\{\overline{\rho_i} \Rightarrow \overline{e_i}\}$: φ then Γ ; Δ , Δ' , Δ'' \vdash case e_c of z: $_{\Delta}\tau$ $\{\overline{\rho_{c_i} \Rightarrow case}\ e_{c_i}$ of w $\{\overline{\rho_i \Rightarrow e_i}\}\}$: φ

(1) Γ ; Δ , Δ' , $\Delta'' \vdash$ case case e_c of $z:_{\Delta}\sigma$ $\{\overline{\rho_{c_i} \Rightarrow e_{c_i}}\}$ of $w:\sigma'$ $\{\overline{\rho_i \Rightarrow e_i}\}: \varphi$

(2) Γ ; Δ , $\Delta' \vdash \mathbf{case} \ e_c \ \mathbf{of} \ z_{:\Delta} \sigma \ \{ \overline{\rho_{c_i} \Rightarrow e_{c_i}} \} : \sigma'$

Proof.

```
(3) \Gamma, w:_{[\Delta,\Delta']}\sigma'; [\Delta,\Delta'], \Delta'' \vdash \rho \Rightarrow e_i:_{[\Delta,\Delta']}^w \varphi
                                                                                                                                                                                                                                  by inv.
Subcase e_c is not in WHNF
(4) \Gamma; \Delta \vdash e_c : \sigma
(5) \overline{\Gamma, z_{[\Delta]}\tau; [\Delta], \Delta' \vdash \rho_{c_i} \Rightarrow e_{c_i} :_{\lceil \Lambda \rceil}^z \sigma'}
                                                                                                                                                                                                                          by inv. (2)
For all alternatives
Subcase \rho_{c_i} = 
(6) \Gamma, z:_{[\Delta]}\tau; [\Delta], \Delta' \vdash e_{c_i} : \sigma'

(7) \Gamma, w:_{[[\Delta],\Delta']}\sigma', \Delta''; [\Delta,\Delta'], \Delta'' \vdash_{alt} \rho \Rightarrow e_i:_{[[\Delta],\Delta']}^{w} \varphi
                                                                                                                                                                                                     by inv. on Alt_{-}(5)
                                                                                                                                                                                                       by irrelevance (3)
Subcase e_{C_i} is not in WHNF
(8) \Gamma, z:_{\lceil \Delta \rceil} \tau; \lceil \Delta \rceil, \Delta', \Delta'' \vdash \mathbf{case} \ e_{c_i} \ \mathbf{of} \ w:_{\lceil \lceil \Delta \rceil, \Delta' \rceil} \sigma' \ \{ \overline{\rho_i \Rightarrow e_i} \} : \varphi
                                                                                                                                                                                        by (6,7) CaseNotWHNF
Subcase e_{c_i} is in WHNF
(8) \Gamma, z:_{\lceil \Delta \rceil} \sigma; [\Delta], \Delta' \Vdash e_{c_i} : \sigma' > \overline{\Delta_i}
                                                                                                                                                                                for some \overline{\Delta_i} in this subcase
(9) \Gamma, w:_{\overline{\Delta_i}} \sigma'; \overline{\Delta_i}, \Delta'' \vdash \rho_j \Longrightarrow e_i :_{\overline{\Delta_i}}^w \varphi
                                                                                                                                                                                                       by irrelevance (3)
                                                                                                                                                                                             for some matching \rho_i
(10) \Gamma, z:_{[\Delta]} \tau; [\Delta], \Delta', \Delta'' \vdash \mathbf{case} \ e_{c_i} \ \mathbf{of} \ w:_{\overline{\Delta_i}} \sigma' \ \{\overline{\rho_i \Rightarrow e_i}\} : \varphi
                                                                                                                                                                                              by (7,8,9) CaseWHNF
Subcase \rho_{c_i} = K \overline{x_{\omega} \tau}
(6) \Gamma, z: \tau, \overline{x}:_{\omega} \overline{\tau}; \cdot, \Delta' \vdash e_{c_i} : \sigma'

(7) \Gamma; w:_{[\Delta']} \sigma'; [\Delta'], \Delta'' \vdash \rho_i \Rightarrow e_i :_{[\Delta']}^w \varphi
                                                                                                                                                                                                     by inv. on Alt 0 (5)
                                                                                                                                                                                                       by irrelevance (3)
Subcase e_{c_i} is not in WHNF
(8) \Gamma, z: \tau, \overline{x:_{\omega}\tau}; \Delta', \Delta'' \vdash \mathbf{case} \ e_{c_i} \ \mathbf{of} \ w:_{[\Delta']}\sigma' \ \{\overline{\rho_i \Rightarrow e_i}\} : \varphi
                                                                                                                                                                                        by (6,7) CaseNotWHNF
Subcase e_{c_i} is in WHNF
(8) \; \Gamma, z :. \tau, \overline{x :_{\omega} \underline{\tau}}; \cdot, \Delta' \Vdash e_{c_i} : \sigma' > \overline{\Delta_i}
                                                                                                                                                                                                         by (6) in subcase
(9) \; \Gamma, w :_{\overline{\Delta_i}} \sigma'; \overline{\Delta_i}, \Delta'' \vdash_{alt} \; \rho_i \Longrightarrow e_i :_{\overline{\Delta_i}}^w \varphi
                                                                                                                                                                                                       by irrelevance (3)
(10)\;\Gamma,z{:}\;\tau,\overline{x{:}_{\omega}\tau};\overline{\Delta_{i}},\Delta^{\prime\prime}\vdash\mathbf{case}\;e_{c_{i}}\;\mathbf{of}\;w_{\;\overline{\Delta_{i}}}\sigma^{\prime}\;\{\overline{\rho_{i}\Longrightarrow e_{i}}\}:\varphi
                                                                                                                                                                                              by (7,8,9) CaseWHNF
Subcase \rho_{c_i} = K \overline{x_{:\omega} \tau y_{:1} \tau}, recalling that e_c is not in WHNF
(6) \Gamma, z:_{[\Delta]} \tau, \overline{x:_{\omega} \tau}, \overline{y:_{[\Delta] \# K_i} \sigma}; [\Delta], \Delta' \vdash e_{c_i} : \sigma'
                                                                                                                                                                               by inv. on Alt N_{\text{Not WHNF}} (5)
(7) \Gamma, w:_{[[\Delta], \Delta']} \sigma'; [[\Delta], \Delta'], \Delta'' \vdash_{alt} \rho_i \Rightarrow e_i :_{[[\Lambda], \Lambda']}^w \varphi
                                                                                                                                                                                                       by irrelevance (3)
```

Proc. ACM Program. Lang., Vol. 10, No. POPL, Article 69. Publication date: January 2026.

Subcase e_{c_i} is not in WHNF

(8)
$$\Gamma$$
, $z:_{[\Delta]} \tau$, $\overline{x:_{\omega} \tau}$, $\overline{y:_{[\Delta] \# K_i} \sigma}$; $[\Delta]$, Δ' , $\Delta'' \vdash \mathbf{case} \ e_{c_i} \ \mathbf{of} \ w:_{[[\Delta], \Delta']} \sigma' \ \{\overline{\rho_i \Rightarrow e_i}\} : \varphi$
Subcase e_{c_i} is in WHNF

(8)
$$\Gamma, z_{:[\Delta]} \tau, \overline{x_{:\omega}} \tau, \overline{y_{:[\Delta] \# K_i}} \overline{\sigma}; [\Delta], \Delta' \Vdash e_{c_i} : \sigma' > \overline{\Delta_i}$$
 by (6) in subcase (9) $\Gamma, w_{:[\Delta], \Delta'} \sigma'; [\Delta], \Delta', \Delta'' \vdash_{alt} \rho_i \Longrightarrow e_i :_{[\Delta], \Delta'}^w \varphi$ by irrelevance (3)

(10)
$$\Gamma$$
, $z:_{\lceil \Delta \rceil} \tau$, $\overline{x}:_{\omega} \tau$, $\overline{y}:_{\lceil \Delta \rceil \# K_i} \overline{\sigma}$; $[\Delta]$, Δ' , $\Delta'' \vdash \mathbf{case} \ e_{c_i} \ \mathbf{of} \ w:_{\lceil \Delta \rceil, \Delta'} \sigma' \ \{\overline{\rho_i \Rightarrow e_i}\} : \varphi$

Commonly to all alternatives subcases:

(11) Γ ; Δ , Δ' , $\Delta'' \vdash \mathbf{case} \ e_c \ \mathbf{of} \ \text{alternatives from (8) or (10)} : \varphi$

Subcase e_c is in WHNF

(4)
$$\Gamma$$
; $\Delta \Vdash e_c : \sigma > \Delta$

$$\begin{array}{ll} \text{(5)}\; \Gamma, z:_{\Delta}\sigma; \Delta, \Delta' \vdash \rho_{c_j} \Rightarrow e_{c_i} :_{\Delta}^z \; \sigma \mapsto \varphi \\ \text{(5)}\; \Gamma, z:_{[D]}\tau; [\Delta] \; , \Delta' \vdash \rho_{c_i} \Rightarrow e_{c_i} :_{[\Delta]}^z \; \sigma \Rrightarrow \varphi \end{array} \qquad \qquad \text{for } \rho_{c_j} \; \text{matches } e_c$$

$$(5) \overline{\Gamma, z_{[D]}\tau; [\Delta], \Delta' \vdash \rho_{c_i} \Rightarrow e_{c_i} :_{[\Delta]}^z \sigma \Rightarrow \varphi}$$

Continue as in the previous subcase, but with Δ instead of $\lceil \Delta \rceil$

C.5 Commuting let

The let floating transformations move lazy let constructs, in and out of other constructs, to further unblock more optimisations. In essence, since let bindings consume resources lazily (by introducing a Δ -variable with usage environment Δ , where Δ is the typing environment of the bound expression), we can intuitively move them around without violating linearity. We prove let floating transformations full-laziness and three local-transformations preserve types and linearity.

THEOREM C.7 (FULL-LAZINESS PRESERVES TYPES).

If Γ ; Δ , $\Delta' \vdash \lambda y$: $_{\pi}\tau$. let x: $_{\Delta}\sigma = e$ in $e' : \varphi$ and y does not occur in ethen Γ ; Δ , $\Delta' \vdash \text{let } x:_{\Delta} \sigma = e \text{ in } \lambda y:_{\pi} \tau$. e'

PROOF.

(1)
$$\Gamma$$
; Δ , $\Delta' \vdash \lambda y$: $_{\pi}\tau$. let x : $_{\Delta}\sigma = e$ in $e' : \sigma' \to \varphi$

Subcase $\pi = 1$

(2)
$$\Gamma$$
; Δ , Δ' , y :₁ $\tau \vdash$ **let** x : $\Delta \sigma = e$ in $e' : \varphi$ by inv. on λI

(3)
$$\Gamma$$
; $\Delta \vdash e : \sigma$ by inv. on (let)

(4)
$$\Gamma, x:_{\Delta}\sigma; \Delta, \Delta', y:_{1}\tau \vdash e' : \varphi$$
 by inv. on (let)
(5) $\Gamma, x:_{\Delta}\sigma; \Delta, \Delta' \vdash \lambda y:_{1}\tau \cdot e' : \sigma' \rightarrow \varphi$ by (λI) (4)

(5)
$$\Gamma, x:_{\Delta}\sigma; \Delta, \Delta' \vdash \lambda y:_{1}\tau. \ e': \sigma' \to \varphi$$
 by (λI) (4)

(6) Γ ; Δ , $\Delta' \vdash \text{let } x: \Lambda \sigma = e \text{ in } \lambda y: 1\tau. e' : \sigma' \rightarrow \varphi \text{ by (let) } (3,5)$ Subcase $\pi = \omega$

As above but x is put in the unrestricted context Γ

THEOREM C.8 (COMMUTING LETS PRESERVE TYPES).

$$1. \quad \Gamma ; \Delta \vdash (\textbf{let} \ v = e \ \textbf{in} \ b) \ a : \varphi \\ \qquad \qquad \Rightarrow \quad \Gamma ; \Delta \vdash \textbf{let} \ v = e \ \textbf{in} \ b \ a : \varphi$$

2.
$$\Gamma$$
; $\Delta \vdash \mathbf{case} \ (\mathbf{let} \ v = e \ \mathbf{in} \ b) \ \mathbf{of} \ \ldots : \varphi \implies \Gamma$; $\Delta \vdash \mathbf{let} \ v = e \ \mathbf{in} \ \mathbf{case} \ b \ \mathbf{of} \ \ldots : \varphi$

3.
$$\Gamma$$
; $\Delta \vdash \text{let } v = (\text{let } v = e \text{ in } b) \text{ in } c : \varphi \implies \Gamma$; $\Delta \vdash \text{let } v = e \text{ in } \text{let } x = b \text{ in } c : \varphi$

1. Commuting Let-app

Proof.

Proc. ACM Program. Lang., Vol. 10, No. POPL, Article 69. Publication date: January 2026.

by Let (3,6)

(1) Γ ; Δ , Δ' , $\Delta'' \vdash (\mathbf{let} \ x :_{\Delta} \sigma = e_1 \ \mathbf{in} \ e_2) \ e_3 : \varphi$	
(2) Γ ; Δ , $\Delta' \vdash \text{let } x:_{\Delta} \sigma = e_1 \text{ in } e_2 : \sigma' \rightarrow_{\pi} \varphi$	by inv. on 1
(3) Γ ; $\Delta'' \vdash e_3 : \sigma'$	by inv. on 1
(4) Γ ; $\Delta \vdash e_1 : \sigma$	by inv. on 2
(5) $\Gamma, x:_{\Delta} \sigma; \Delta, \Delta' \vdash e_2 : \sigma' \rightarrow_{\pi} \varphi$	by inv. on 2
(6) Γ , $x:_{\Delta}\sigma$; Δ , Δ' , $\Delta'' \vdash e_2 e_3 : \varphi$	by $\lambda_{\pi}E$
(7) Γ ; Δ , Δ' , $\Delta'' \vdash \mathbf{let} \ x :_{\Delta} \sigma = e_1 \ \mathbf{in} \ e_2 \ e_3 : \varphi$	by let (4,6)

2. Commuting let-case

Proof.

```
(2) \Gamma; \Delta, \Delta' \vdash \mathbf{let} \ x :_{\Delta} \sigma = e_1 \ \mathbf{in} \ e_2 : \sigma'
                                                                                                                                                                                                by inv. on 1
(3) \Gamma; \Delta \vdash e_1 : \sigma
                                                                                                                                                                                                by inv. on 2
(4) \Gamma, x:_{\Lambda}\sigma; \Delta, \Delta' \vdash e_2 : \sigma'
                                                                                                                                                                                                by inv. on 2
Subcase e_2 is in WHNF
(5) \ \overline{\Gamma, z_{:\Delta,\Delta'}\sigma'; \Delta, \Delta', \Delta''} \vdash_{alt} \rho \Rightarrow e_3 :_{\Delta,\Delta'}^z \sigma' \mapsto \varphi
                                                                                                                                                                                                by inv. on 1
(6) \Gamma, x:_{\Delta} \sigma; \Delta, \Delta', \Delta'' \vdash \mathbf{case} \ e_2 \ \mathbf{of} \ z:_{\Delta, \Delta'} \sigma' \ \{ \overline{\rho \Rightarrow e_3} \} : \varphi
                                                                                                                                                                             by CaseWHNF (4,5)
(7) \Gamma; \Delta, \Delta', \Delta'' \vdash \text{let } x:_{\Delta} \sigma = e_1 \text{ in case } e_2 \text{ of } z:_{\Delta,\Delta'} \sigma' \{ \overline{\rho \Rightarrow e_3} \} : \varphi
                                                                                                                                                                                                 by Let (3,6)
Subcase e_2 is not in WHNF
(5) \Gamma, z:_{[\Delta, \Delta']} \sigma'; [\Delta, \Delta'], \Delta'' \vdash_{alt} \rho \Rightarrow e_3 :_{[\Delta, \Delta']}^z \sigma' \Rightarrow \varphi
                                                                                                                                                                                                by inv. on 1
(6) \Gamma, x:_{\Delta} \sigma; \Delta, \Delta', \Delta'' \vdash \mathbf{case} \ e_2 \ \mathbf{of} \ z:_{\Delta, \Delta'} \sigma' \{ \overbrace{\rho \Rightarrow e_3} \} : \varphi
                                                                                                                                                                             by CaseWHNF (4,5)
```

(1) Γ ; Δ , Δ' , $\Delta'' \vdash \mathbf{case let} \ x:_{\Delta} \sigma = e_1 \ \mathbf{in} \ e_2 \ \mathbf{of} \ z:_{\Delta,\Delta'} \sigma' \ \{ \overline{\rho \Rightarrow e_3} \} : \varphi$

(7) Γ ; Δ , Δ' , $\Delta'' \vdash \text{let } x:_{\Delta} \sigma = e_1 \text{ in case } e_2 \text{ of } z:_{\Delta,\Delta'} \sigma' \{ \overline{\rho \Rightarrow e_3} \} : \varphi$

3. Commuting let-let

Proof.

 $(1) \ \Gamma, \Delta; \Delta', \Delta'' \vdash \mathbf{let} \ x:_{\Delta,\Delta'} \sigma' = (\mathbf{let} \ y:_{\Delta} \sigma' = e_1 \ \mathbf{in} \ e_2) \ \mathbf{in} \ e_3 : \varphi$ $(2) \ \Gamma; \Delta, \Delta' \vdash \mathbf{let} \ y:_{\Delta} \sigma' = e_1 \ \mathbf{in} \ e_2 : \sigma'$ $(3) \ \Gamma, x:_{\Delta,\Delta'} \sigma'; \Delta, \Delta', \Delta'' \vdash e_3$ $(4) \ \Gamma; \Delta \vdash e_1 : \sigma$ $(5) \ \Gamma, y:_{\Delta} \sigma'; \Delta, \Delta' \vdash e_2 : \sigma'$ $(6) \ \Gamma, y:_{\Delta} \sigma'; \Delta, \Delta', \Delta'' \vdash \mathbf{let} \ x:_{\Delta,\Delta'} \sigma' = e_2 \ \mathbf{in} \ e_3 : \varphi$ $(7) \ \Gamma; \Delta, \Delta', \Delta'' \vdash \mathbf{let} \ y:_{\Delta} \sigma' = e_1 \ \mathbf{in} \ \mathbf{let} \ x:_{\Delta,\Delta'} \sigma' = e_2 \ \mathbf{in} \ e_3 : \varphi$ by Let (3,5) and Weaken by Let (4,6)

Theorem C.9. If Γ ; $\Delta_1, \Delta_2, \Delta_3 \vdash \mathbf{case}\ e\ \mathbf{of}\ \{\overline{\rho_j \Rightarrow e_j}, \rho \Rightarrow E[e_1], \overline{\rho_k \Rightarrow e_k}\}: \varphi\ and\ e_1\ does\ not\ use\ any\ pattern\ variables\ introduced\ by\ \rho\ , the\ case\ binder\ nor\ variables\ bound\ in\ context\ E[-]\ then\ <math>\Gamma$; $\Delta_1, \Delta_2, \Delta_3 \vdash \mathbf{let}\ x = e_1\ \mathbf{in}\ \mathbf{case}\ e\ \mathbf{of}\ \{\overline{\rho_j \Rightarrow e_j}, \rho \Rightarrow E[x], \overline{\rho_k \Rightarrow e_k}\}: \varphi\ , for\ some\ fresh\ x.$ PROOF.

Subcase: e is in WHNF

(1) Γ ; $\Delta_1 \Vdash e : \sigma > \overline{\Delta_i}$ by inversion

```
(2) \Gamma; \Delta_1 \vdash e : \sigma
                                                                                                                                        by inversion and constructor rule if needed
(3) \Gamma, z:_{\overline{\Delta_{i}}}\tau; \overline{\Delta_{i}}, \Delta_{2}, \Delta_{3} \vdash_{alt} \rho_{j} \Rightarrow e':_{\overline{\Delta_{i}}}z \Rightarrow \varphi

(4) \Gamma, z:_{[\Delta_{1}]}\sigma; [\Delta_{1}], \Delta_{2}, \Delta_{3} \vdash_{alt} \rho \Rightarrow e':_{[\Delta_{1}]}z \Rightarrow \varphi
                                                                                                                                                                                                                      by inversion
                                                                                                                                                                                                                      by inversion
(5) e matches \rho_i
                                                                                                                                                                                                                      by inversion
Subsubcase: \rho_i corresponds to branch containing E[e_1]
Subsubsubcase: (3) derived by AltN<sub>WHNF</sub>
\Gamma, z:_{\overline{\Delta_i}} \sigma, \overline{x:_{\omega} \tau}, \overline{y_i:_{\Delta_i} \sigma_i}^n; \overline{\Delta_i}, \Delta_2, \Delta_3 \vdash E[e_1]: \varphi
                                                                                                                                                                                                                      by inversion
\Gamma, z: \overline{\Lambda_i} \sigma, \overline{x:_{\omega} \tau}, \overline{y_i:_{\Delta_i} \sigma_i}^n; \Delta', \Delta_2 \vdash e_1 : \sigma'
                                                                                                                                                                    by inversion, with \Delta' \subseteq \Delta_i \subseteq \Delta_1
\Gamma, x:_{\Delta', \Delta_2} \sigma', z:_{\overline{\Lambda}} \sigma, \overline{x:_{\omega} \tau}, \overline{y_i:_{\Delta_i}, \sigma_i}^n; \Delta', \Delta_2 \vdash x: \sigma'
                                                                                                                                                                                               by Var_{\Delta}, with x fresh
\Gamma; \Delta', \Delta_2 \vdash e_1 : \sigma'
                                                                                                                                                                                                                by assumption
\Gamma, x :_{\Delta', \Delta_2} \sigma', z :_{\overline{\Delta_i}} \sigma, \overline{x :_{\omega} \tau}, \overline{y_i :_{\Delta_i} \sigma_i}^n; \overline{\Delta_i}, \Delta_2, \Delta_3 \vdash E[x] : \varphi
                                                                                                                                                                                         by context instantiation
\Gamma, x:_{\Delta', \Delta_2} \sigma', z:_{\overline{\Delta_i}} \tau; \overline{\Delta_i}, \Delta_2, \Delta_3 \vdash_{alt} \rho_j \Rightarrow E[x] :_{\overline{\Delta_i}}^z \sigma \mapsto \varphi
                                                                                                                                                                                                                    by AltN<sub>WHNF</sub>
\Gamma, x:_{\Delta', \Delta_2} \sigma'; \Delta_1, \Delta_2, \Delta_3 \vdash \mathbf{case} \ e \ \mathbf{of} \ \{ \overline{\rho_j \Rightarrow e_j}, \rho \xrightarrow{\cdot} E[x], \overline{\rho_k \Rightarrow e_k} \} : \varphi
                                                                                                                                                                                                                    by Case<sub>WHNF</sub>
\Gamma; \Delta_1, \Delta_2, \Delta_3 \vdash \text{let } x:_{\Delta', \Delta_2} \tau = e_1 \text{ in case } e \text{ of } \{\overline{\rho_j \Rightarrow e_j}, \rho \Rightarrow E[x], \overline{\rho_k \Rightarrow e_k}\} : \varphi
                                                                                                                                                                                                                                     by Let
Subsubsubcase: (3) derived by Alt0
\Gamma, z:_{\overline{\Delta_i}} \sigma\left[\cdot/\overline{\Delta_i}\right]_z, \overline{x:_{\omega}\tau}; \overline{\Delta_i}, \Delta_2, \Delta_3\left[\cdot/\overline{\Delta_i}\right] \vdash E[e_1]: \varphi
                                                                                                                                                                                                                     by inversion
\Gamma, z :: \sigma, \overline{x :_{\omega} \tau}; \Delta_2 \vdash e_1 : \sigma'
                                                                                                                                                                                                                      by inversion
\Gamma; \Delta_2 \vdash e_1 : \sigma'
                                                                                                                                                                                                                by assumption
\Gamma, x:_{\Delta_2} \sigma', z:.\sigma, \overline{x:_{\omega} \tau}; \Delta_2 \vdash x:\sigma'
                                                                                                                                                                                               by Var_{\Delta}, with x fresh
\Gamma, x:_{\Delta_2} \sigma', z:_{\overline{\Lambda_i}} \sigma \left[ \cdot / \overline{\Delta_i} \right]_z, \overline{x:_{\omega} \tau}; \overline{\Delta_i}, \Delta_2, \Delta_3 \left[ \cdot / \overline{\Delta_i} \right] + E[x] : \varphi
                                                                                                                                                                                         by context instantiation
\Gamma, x:_{\Delta_2} \sigma'; \Delta_1, \Delta_2, \Delta_3 \vdash \mathbf{case} \ e \ \mathbf{of} \ \{\overline{\rho_j \Rightarrow e_j}, \rho \Rightarrow E[x], \overline{\rho_k \Rightarrow e_k}\} : \varphi
                                                                                                                                                                                             by Alt0 and Case<sub>WHNF</sub>
\Gamma; \Delta_1, \Delta_2, \Delta_3 \vdash \text{let } x:_{\Delta_2} \tau = e_1 \text{ in case } e \text{ of } \{\overline{\rho_j \Rightarrow e_j}, \rho \Rightarrow E[x], \overline{\rho_k \Rightarrow e_k}\} : \varphi
                                                                                                                                                                                                                                     by Let
Subsubsubcase: (3) derived by Alt_
\Gamma, z:_{\overline{\Lambda_i}}\sigma; \Delta_2, \Delta_3 \vdash E[e_1] : \varphi
                                                                                                                                                                                                                     by inversion
\Gamma, z:_{\overline{\Delta_i}}\sigma; \Delta_2 \vdash e_1 : \sigma'
                                                                                                                                                                                                                     by inversion
\Gamma, z: \overline{\Lambda}, \sigma, x: \overline{\Lambda} \sigma'; \Delta_2 \vdash x : \sigma'
                                                                                                                                                                                                by Var_{\Delta}, with x fresh
\Gamma, z:_{\overline{\Lambda_i}}\sigma, x:_{\Delta_2}\sigma'; \Delta_2, \Delta_3 \vdash E[e_1] : \varphi
                                                                                                                                                                                         by context instantiation
\Gamma, x:_{\Delta_2} \sigma'; \Delta_1, \Delta_2, \Delta_3 \vdash \mathbf{case} \ e \ \mathbf{of} \ \{\overline{\rho_i \Rightarrow e_i}, \rho \Rightarrow E[x], \overline{\rho_k \Rightarrow e_k}\} : \varphi
                                                                                                                                                                                            by Alt_ and CaseWHNF
\Gamma; \Delta_1, \Delta_2, \Delta_3 \vdash \text{let } x:_{\Delta_2} \tau = e_1 \text{ in case } e \text{ of } \{\overline{\rho_j \Rightarrow e_j}, \rho \Rightarrow E[x], \overline{\rho_k \Rightarrow e_k}\} : \varphi
                                                                                                                                                                                                                                     by Let
Subsubcase: \rho_i does not correspond to branch containing E[e_1]
(6) \Gamma, z:_{[\Delta_1]}\sigma; [\Delta_1], \Delta_2, \Delta_3 \vdash_{alt} \rho \Rightarrow E[e_1]:_{[\Lambda_1]}^z \sigma \Rightarrow \varphi
                                                                                                                                                                                                              this subsubcase
Subsubsubcase: (6) derived by AltN<sub>Not WHNF</sub>
\Gamma, z:_{[\Delta_1]} \sigma, \overline{x:_{\omega} \tau}, \overline{y_i:_{\Delta_i} \sigma_i}; [\Delta_1], \Delta_2, \Delta_3 \vdash E[e_1] : \varphi \text{ and } \overline{\Delta_i} = \overline{\Delta_1 \# K_i}^n
                                                                                                                                                                                                                      by inversion
\Gamma, z:_{[\Delta_1]}\sigma, \overline{x:_{\omega}\tau}, \overline{y_i:_{\Delta_i}\sigma_i}; \Delta_2 \vdash e_1 : \sigma'
                                                                                                                                                                                                                      by inversion
\Gamma, x:_{\Delta_2} \sigma', z:_{[\Delta_1]} \sigma, \overline{x:_{\omega} \tau}, \overline{y_i:_{\Delta_i} \sigma_i}; \Delta_2 \vdash x: \sigma'
                                                                                                                                                                                               by Var_{\Delta}, with x fresh
\Gamma, x:_{\Delta_2} \sigma', z:_{[\Delta_1]} \sigma, \overline{x:_{\omega} \tau}, \overline{y_i:_{\Delta_i} \sigma_i}; [\Delta_1], \Delta_2, \Delta_3 \vdash E[x] : \varphi \text{ and } \overline{\Delta_i} = \overline{\Delta_1 \# K_i}^n
                                                                                                                                                                                        by context instantiation
\Gamma, x:_{\Delta_2} \sigma'; \Delta_1, \Delta_2, \Delta_3 \vdash \mathbf{case} \ e \ of \ \{\overline{\rho_i \Rightarrow e_i}, \rho \Rightarrow E[x], \overline{\rho_k \Rightarrow e_k}\} : \varphi \ \text{by AltN}_{\text{Not WHNF}} \ \text{and Case}_{\text{WHNF}}
\Gamma; \Delta_1, \Delta_2, \Delta_3 \vdash \text{let } x:_{\Delta_2} \tau = e_1 \text{ in case } e \text{ of } \{\overline{\rho_i \Rightarrow e_i}, \rho \Rightarrow E[x], \overline{\rho_k \Rightarrow e_k}\} : \varphi
                                                                                                                                                                                                                                     by Let
Subsubsubcase: (6) derived by Alt_ or Alt0
Identical to subsubsubcases above.
Subcase: e is not in WHNF
(1) \Gamma; \Delta_1 \vdash e : \sigma
                                                                                                                                                                                                                     by inversion
(2) \Gamma, z:_{[\Delta_1]} \sigma; [\Delta_1], \Delta_2, \Delta_3 \vdash_{alt} \rho \Rightarrow e' :_{[\Delta_1]}^z \sigma \Rightarrow \varphi
                                                                                                                                                                                                                      by inversion
(3) \Gamma, z:_{[\Delta_1]}\sigma; [\Delta_1], \Delta_2, \Delta_3 \vdash_{alt} \rho \Rightarrow E[e_1]:_{[\Delta_1]}^z \sigma \Rightarrow \varphi
                                                                                                                                                                                                                      by inversion
```

Subsubcase: (3) derived by AltN_{Not WHNF}, Alt_ or Alt0

Identical to Subsubsubcases above.

C.6 η -conversions

The η -conversion transformations are η -expansion and η -reduction. In both transformations, linearity is preserved since the resources used to type the function f do not change neither when the lambda and its argument are removed, nor when we add a lambda and apply f to the bound argument.

Theorem C.10 (η -expansion preserves types). If Γ ; $\Delta \vdash f : \sigma \rightarrow_{\pi} \varphi$ then Γ ; $\Delta \vdash \lambda x$. $f : x : \sigma \rightarrow_{\pi} \varphi$

Subcase f is linear

- (1) Γ ; $\Delta \vdash f : \sigma \multimap \varphi$
- (2) Γ ; x: $_1\tau \vdash x : \sigma$

(3)
$$\Gamma$$
; Δ , x : $_1\tau \vdash f x : \varphi$

by λE by λI

$$(4) \; \Gamma; \Delta \vdash (\lambda x :_1 \tau. \; f \; x) : \sigma \multimap \varphi$$

Subcase f is unrestricted As above but x is introduced in Γ and functions are unrestricted

Theorem C.11 (η -reduction preserves types). If Γ ; $\Delta \vdash \lambda x$. $f x : \sigma \rightarrow_{\pi} \varphi$ then Γ ; $\Delta \vdash f : \sigma \rightarrow_{\pi} \varphi$ Proof.

(1) Γ ; $\Delta \vdash (\lambda x :_{\pi} \tau. f x) : \sigma \rightarrow_{\pi} \varphi$ Subcase $\pi = 1$

(2) Γ ; Δ , x:₁ $\tau \vdash f x : \varphi$ by inv. on λI by inv. on λE

(3) Γ ; $\Delta \vdash f : \sigma \rightarrow_1 \varphi$

Subcase $\pi = \omega$

As above but x is introduced in Γ

C.7 Binder Swap

The binder swap transformation applies to case expressions whose scrutinee is a single variable x, and it substitutes occurrences of x in the case alternatives for the case binder z. If x is a linear resource, x cannot occur in the case alternatives (as we conservatively consider variables are not in WHNF), so the substitution preserves types vacuously. Otherwise, x can be freely substituted by z, since z is also an unrestricted resource (it's usage environment is empty because x is unrestricted).

Theorem C.12 (Binder-Swap Preserves Types). If Γ ; $\Delta \vdash \mathbf{case} \ x \ \mathbf{of} \ z \ \{\overline{\rho_i \Rightarrow e_i}\} : \varphi \ then \ \Gamma$; $\Delta \vdash \mathbf{case} \ x \ \mathbf{of} \ z \ \{\overline{\rho_i \Rightarrow e_i[z/x]}\} : \varphi$ Proof.

Subcase *x* is linear

- (1) Γ ; Δ , $x:_1\tau \vdash \mathbf{case} \ x \ \mathbf{of} \ z:_{[x]}\tau \ \{\overline{\rho \Rightarrow e}\} : \varphi$
- by inv. on Case_{Not WHNF} (2) Γ ; x: $_1\tau \vdash x : \sigma$

 $(3) \ \overline{\Gamma, z_{:[x]}\tau; \Delta, [x_{:1}\tau] \vdash_{alt} \rho \Rightarrow e :^{z}_{[x]} \sigma \Rightarrow \varphi}$ by inv. on Case_{Not WHNF}

Proc. ACM Program. Lang., Vol. 10, No. POPL, Article 69. Publication date: January 2026.

$$(4) \ \overline{\Gamma, z:_{[x]} \tau; \Delta, [x:_1 \tau] \vdash_{alt} \rho \Rightarrow e[z/x] :_{[x]}^z \sigma \Rrightarrow \varphi}$$

(5) Γ ; Δ , x:₁ τ \vdash **case** x **of** z:[x] τ { $\overline{\rho \Rightarrow e[z/x]}$ } : φ Subcase x is unrestricted

- (1) $\Gamma, x:_{\omega} \tau; \Delta \vdash \mathbf{case} \ x \ \mathbf{of} \ z:_{\tau} \{ \overline{\rho \Rightarrow e} \} : \varphi$
- (2) Γ , $x:_{\omega}\tau$; $\cdot \vdash x : \sigma$
- $(3) \ \overline{\Gamma, x:_{\omega}\tau, z:_{\tau}; \Delta \vdash_{alt} \rho \Rightarrow e:_{\tau}^{z} \sigma \Rrightarrow \varphi}$
- (4) Γ , z:. τ ; · \vdash z: σ
- (5) $\overline{\Gamma, z: \tau; \Delta \vdash_{alt} \rho \Rightarrow e[z/x] : \sigma \Rightarrow \varphi}$
- (6) $\Gamma, x:_{\omega} \tau; \Delta \vdash \mathbf{case} \ x \ \mathbf{of} \ z:_{\varepsilon} \tau \ \{ \overline{\rho \Rightarrow e[z/x]} \} : \varphi$

by x cannot occur in e be it's proof irrelevant by $Case_{\text{Not WHNF}}$

by inv. on $Case_{\mathrm{Not}}$ Whnf by inv. on $Case_{\mathrm{Not}}$ Whnf by Var_{Δ} by unr. subst. lemma (3,4) by $Weaken_{\omega}$ and $Case_{\mathrm{Not}}$ Whnf

Received 2025-07-10; accepted 2025-11-06