

Safe Zero-cost Coercions for Haskell

JOACHIM BREITNER

Karlsruhe Institute of Technology
(e-mail: breitner@kit.edu)

RICHARD A. EISENBERG

University of Pennsylvania
(e-mail: eir@cis.upenn.edu)

SIMON PEYTON JONES

Microsoft Research
(e-mail: simonpj@microsoft.com)

STEPHANIE WEIRICH

University of Pennsylvania
(e-mail: sweirich@cis.upenn.edu)

Abstract

Generative type abstractions – present in Haskell, OCaml, and other languages – are useful concepts to help prevent programmer errors. They serve to create new types that are distinct at compile time but share a run-time representation with some base type. We present a new mechanism that allows for zero-cost conversions between generative type abstractions and their representations, even when such types are deeply nested. We prove type safety in the presence of these conversions and have implemented our work in GHC.

1 Introduction

Modular languages support *generative type abstraction*, the ability for programmers to define application-specific types, and rely on the type system to distinguish between these new types and their underlying representations. Type abstraction is a powerful tool for programmers, enabling both flexibility (implementors can change representations) and security (implementors can maintain invariants about representations). Typed languages provide these mechanisms with zero run-time cost – there should be no performance penalty for creating abstractions – using mechanisms such as ML’s module system (Milner *et al.*, 1997) and Haskell’s **newtype** declaration (Marlow, 2010).

For example, a Haskell programmer might create an abstract type for HTML data, representing them as Strings (Figure 1). Although String values use the same patterns of bits in memory as HTML values, the two types are distinct. That is, a String will not be accepted by a function expecting an HTML. The data constructor `MkHTML` converts a String to an HTML (see function text), while using `MkHTML` in a pattern converts in the other direction (see function `unHTML`). By exporting

```

module Html( HTML, text, unHTML, ... ) where

newtype HTML = MkHTML String

unHTML :: HTML → String
unHTML (MkHTML s) = s

text :: String → HTML
text s = MkHTML (escapeSpecialCharacters s)

```

Figure 1. An abstraction for HTML values

the type `HTML`, but not its data constructor, module `Html` ensures that the type `HTML` is *abstract* – clients cannot make arbitrary strings into `HTML` – and thereby prevent cross-site scripting attacks.

Using **newtypes** for abstraction in Haskell has always suffered from an embarrassing difficulty. Suppose that in the module `Html`, the programmer wants to break `HTML` data into a list of lines, using the standard Haskell library function `lines :: String → [String]`:

```

linesH :: HTML → [HTML]
linesH h = map MkHTML (lines (unHTML h))

```

To get the resulting `[HTML]` we are forced to map `MkHTML` over the list. Operationally, this map is the identity function – the run-time representation of `[String]` is identical to `[HTML]` – *but it will carry a run-time cost nevertheless*. The optimiser in the Glasgow Haskell Compiler (GHC) is powerless to fix the problem because it works over a *typed* intermediate language; the `MkHTML` constructor changes the type of its operand, and hence cannot be optimised away. There is nothing that the programmer can do to prevent this run-time cost. What has become of the claim of zero-overhead abstraction?

In this paper we describe a robust, simple mechanism that programmers can use to solve this problem, making the following contributions:

- We describe the design of *safe coercions* (Section 2), which introduces the function

```
coerce :: Coercible a b ⇒ a → b
```

and the new constraint `Coercible`. This function performs a zero-cost conversion between two types `a` and `b` that have the same representation. The crucial question becomes *for which types is the `Coercible` constraint satisfiable?* We describe how the constraint can be formed and used in Section 2.

- We formalise `Coercible` by translation into GHC’s intermediate language System FC, augmented with the concept of *roles* (Section 2.2), adapted from prior work (Weirich et al., 2011). One new contribution of this work is a simplification of the roles idea; we formalise this simpler system and give the usual statements of preservation and progress in Section 4.

- The appendices contain a complete proof of type safety and type erasure for the variant of System FC described in this article. This proof is the most detailed version of a proof of the safety System FC we are aware of; it serves well as a template for proofs about extensions to System FC. Although the proof broadly echoes prior work, we present it as a novel contribution in its level of detail.
- Adding safe coercions to the source language raises new issues for abstract types, and for the coherence of type elaboration. We articulate the issues, and introduce *role annotations* to solve them (Section 3).
- It would be too onerous to insist on programmer-supplied role annotations for every type, so we give a *role inference algorithm* in Section 4.6.
- The precise algorithm used to simplify and solve Coercible constraints is subtle. It appears in Section 5.
- To support our claim of practical utility, we have implemented the whole scheme in GHC (Section 6), and evaluated it against thousands of Haskell libraries (Section 6.5). Our work also finally resolves a notorious and long-standing bug in GHC (#1496), which concerns the interaction of newtype coercions with type families (Section 6.1).

We build on earlier work on roles (Weirich *et al.*, 2011), which offered a very expressive, but very complicated, system of roles. In this paper we find a sweet spot offering a considerably simpler system in exchange for a minor loss of expressiveness. This article is a revised and expanded version of our ICFP'14 paper (Breitner *et al.*, 2014a), and describes the implementation as it has been refined since the original publication (Section 5).

As this work demonstrates, the interactions between type abstraction and advanced type system features, such as type families and GADTs, are subtle. Although we focus on Haskell and System FC here, the ability to create and enforce zero-cost type abstraction is not unique to Haskell – notably the ML module system also provides this capability, and more. As a result, OCaml developers are now grappling with similar difficulties. We discuss the connection between roles and OCaml's variance annotations (Section 7), as well as other related work.

2 The design and interface of Coercible

We begin by focusing exclusively on the programmer's-eye-view of safe coercions. His entry point to the story is the function

```
coerce :: Coercible a b => a -> b
```

that allows him to convert values between two types a and b . This is reminiscent of the infamous function `unsafeCoerce :: a -> b`, which likewise converts between two types, with the crucial difference that `coerce` works only if it is *safe* to do so, i.e., when the compiler can determine that the conversion will not compromise type safety. This relation between the type a and b is expressed by the new primitive constraint `Coercible a b`. As a constraint, it looks like a type class, but shares only the syntax and the name space with them.

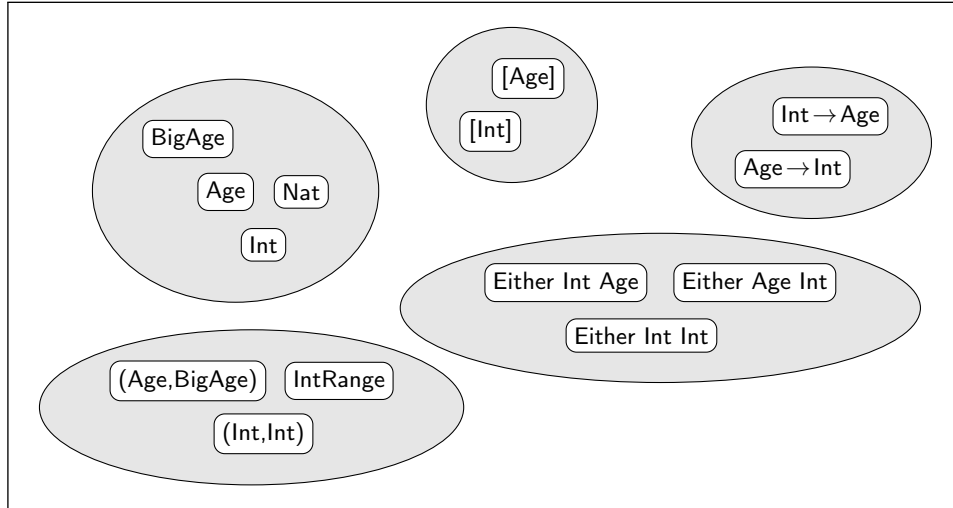


Figure 2. Coercible relates types with identical run-time representation

The key principle is this: *If two types σ and τ are related by `Coercible σ τ`* , then σ and τ have bit-for-bit identical run-time representations. Moreover, as you can see from the type of `coerce`, if `Coercible σ τ` holds then `coerce` can convert a value of type σ to one of type τ with no runtime cost. And that’s it!

The crucial question, to which we devote the rest of this section, becomes this: exactly when does `Coercible σ τ` hold? To whet your appetite consider these declarations:

```

newtype Age      = MkAge Int
newtype BigAge   = MkBig Age
newtype Nat      = MkNat Int
newtype IntRange = MkIR (Int,Int)
    
```

Here are some coercions that hold, so that a single call to `coerce` suffices to convert between the two types:

- `Coercible Int Age`: We can coerce from `Int` to `Age` at zero cost, as this corresponds to simply using the `MkAge` constructor,
- `Coercible Age Int`: and the reverse, as if we were pattern matching on `MkAge`.
- `Coercible BigAge Int`: We can unwrap two steps at once,
- `Coercible BigAge Nat`: and coerce between different newtypes that happen to have the same representation.
- `Coercible [Age] [Int]`: We can lift coercions over lists,
- `Coercible (Either Int Age) (Either Int Int)`: and over `Either`.
- `Coercible (Either Int Age) (Either Age Int)`: It also works if the first argument of `Either` must be coerced in one direction, and the second in the other.
- `Coercible (Int \rightarrow Age) (Age \rightarrow Int)`: All this works over function arrows too.
- `Coercible (Age, BigAge) IntRange`: And even quite complex coercions like this are handled with one call to `coerce`.

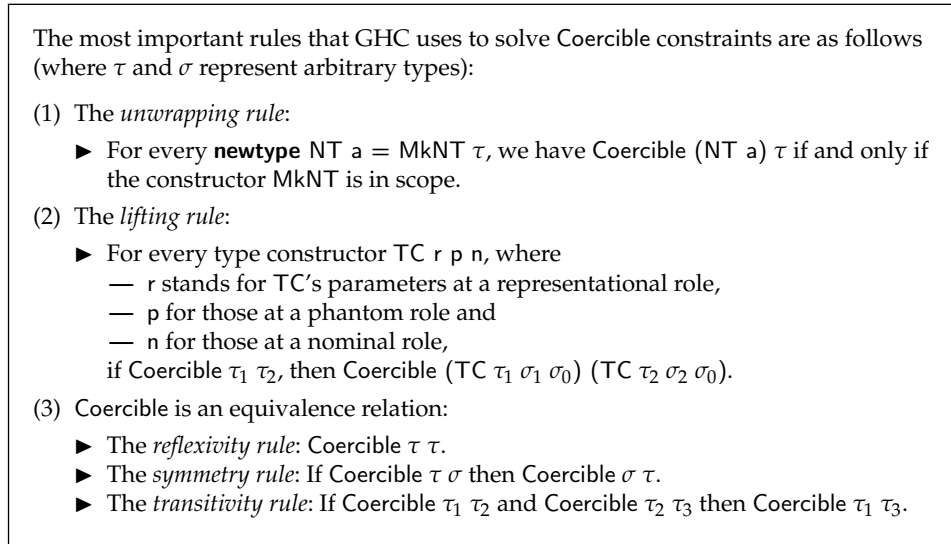


Figure 3. Coercible formation rules (pragmatic summary)

Figure 2 visualises these coercions and shows that Coercible is constructed to be an equivalence relation: It partitions all Haskell types into equivalence classes, so that in each such class, every type has the same run-time representation, and one can convert between any two types that are in the same group, in either direction, with a single call to coerce; but not between types of different groups.

The rest of this section describes the basic rules for determining when one type is Coercible to another; see Figure 3 for a concise summary. The algorithm used in GHC to actually solve Coercible constraints is described in detail in Section 5. Figure 5 contains the full list of coercion formation rules, albeit in the language of System FC instead of Haskell.

2.1 Coercing newtypes

We expect Coercible to relate a newtype with its base type; this is the most obvious rule for Coercible. In our example, this solves the following constraints:

- $\text{Coercible } \text{Int } \text{Age}$
- $\text{Coercible } \text{Age } \text{BigAge}$
- $\text{Coercible } (\text{Int}, \text{Int}) \text{IntRange}$

Notice that each of these rules unwraps just one layer of the newtype, so we call them the *unwrapping rules*.

The newtype-unwrapping rules (i.e., (1) in Figure 3) are available *only if the corresponding newtype data constructor* (e.g. MkNT) *is in scope*; this is required to preserve abstraction, as we explain in Section 3.1.

2.2 Type constructors and roles

The type that we want to coerce might appear as an argument to a type constructor. A type constructor is what forms a parameterised type, and could be a data type, newtype, the function type, or a built-in data type like a tuple or IO. To coerce a composite type, each type constructor has its *lifting rule*, as shown in Figure 3, which lifts coercions through the type constructor.

The shape of the lifting rule depends on the so-called *roles* of the type constructor's parameters. Each type parameter of a type constructor has one of the three possible roles: *representational*, *phantom* and *nominal*. The following subsections explain the meaning of these roles, and how the roles ensure that lifting rules do not give rise to coercions between types with different run-time representations, which would violate type safety.

A role annotation such as

type role Either **representational representational**

is used to assign the roles. Once defined, the roles of a type constructor are the same in every scope, regardless of whether the concrete definition of that type is available in that scope. The compiler checks if these assignments are compatible with how the parameter is used in the definition of the type constructor (Section 4.3). Without a role annotation, a role inference algorithm (Section 4.6) calculates the most liberal allowed role assignment (which, we prove, must be unique), so in practice, role annotations are rarely needed and only used in special circumstances (Section 3.1). We write out some superfluous role annotations to aid the reader.

2.3 Coercing representational type parameters

The most common role is *representational*. It is the role that is assigned to the type parameters of ordinary newtypes and data types that do not discern the actual choice of the type parameter. For example:

type role Maybe **representational**

type role [] **representational**

type role Either **representational representational**

The Coercible rule for these type constructors are:

- ▶ If Coercible $\tau \sigma$ then Coercible (Maybe τ) (Maybe σ).
- ▶ If Coercible $\tau \sigma$ then Coercible $[\tau]$ $[\sigma]$.
- ▶ If Coercible $\tau_1 \sigma_1$ and Coercible $\tau_2 \sigma_2$ then Coercible (Either $\tau_1 \tau_2$) (Either $\sigma_1 \sigma_2$).

These rules are just as you would expect: for example, the type Maybe τ and Maybe σ have the same run-time representation if and only if τ and σ have the same representation.

Most primitive type constructors also have representational roles for their arguments. For example, the domain and co-domain of arrow types are representational, as if we had

type role (\rightarrow) **representational representational**,

giving rise to the following Coercible rule:

- ▶ If Coercible $\tau_1 \sigma_1$ and Coercible $\tau_2 \sigma_2$ then Coercible $(\tau_1 \rightarrow \tau_2) (\sigma_1 \rightarrow \sigma_2)$.

As another example, the type `IORef` has a representational parameter, so expressions of type `IORef Int` can be converted to type `IORef Age` for zero cost (and outside of the IO monad).

Returning to the introduction, we can now write `linesH` very directly, thus:

```
linesH :: HTML → [HTML]
linesH = coerce lines
```

In this case, the call to `coerce` gives rise to a constraint `Coercible (String → [String]) (HTML → [HTML])`, which gets simplified to `Coercible String HTML` using the lifting rules for arrow and list types, and then solved by the unwrapping rule for the newtype `HTML`.

2.4 Coercing phantom type parameters

A type parameter has a *phantom* role if it does not occur at all in the definition of the type, or if it does, then only as a phantom parameter of another type constructor. For example, these declarations

```
data Phantom b = Phantom
data NestedPhantom b = MkNP [Phantom b] | SomethingElse
```

both have parameter `b` at a phantom role:

```
type role Phantom phantom
type role NestedPhantom phantom
```

When do the types `Phantom τ` and `Phantom σ` have the same run-time representation? Always! Therefore, we have the rules

- ▶ Coercible `(Phantom τ) (Phantom σ)`
- ▶ Coercible `(NestedPhantom τ) (NestedPhantom σ)`

and `coerce` can be used to change the phantom parameter arbitrarily.

Such a change can defeat the purpose of the phantom type, but that is an issue of abstraction, not of type safety, and addressed in Section 3.1.

2.5 Coercing nominal type parameters

In contrast, the *nominal* role induces the strictest preconditions for Coercible rules. This role is assigned to a parameter that possibly affects the run-time representation of a type, commonly because it is passed to a type family. For example, consider the following code

```
type family EncData a where
  EncData String = (ByteString, Encoding)
  EncData HTML = ByteString
type role EncData nominal
```

```
data Encoding = ...
data EncText a = MkET (EncData a)
type role EncText nominal
```

Even though we have Coercible HTML String, it would be wrong to accept the constraint Coercible (EncText HTML) (EncText String), because these two types have quite different run-time representations! Therefore, there are no rules that change a nominal parameter of a type constructor.

All parameters of a type *family*¹ have nominal role, because they could be inspected by the type family instances. For similar reasons, the non-uniform parameters to GADTs are also required to be nominal. Type classes also use nominal role for their type parameters; see Section 3.2.

2.6 Coercing multiple type parameters

A type constructor can have multiple type parameters, each at a different role. In that case, an appropriate constraint for each type parameter is used:

```
data Params r p n = Con1 (Maybe r) | Con2 (EncData n)
type role Params representational phantom nominal
```

Hence, following (2) in Figure 3, we get:

- ▶ If Coercible $\tau_1 \tau_2$ then Coercible (Params $\tau_1 \sigma_1 \sigma_0$) (Params $\tau_2 \sigma_2 \sigma_0$).

2.7 Inverting the lifting rule

For a data type constructor such as Maybe, there is only one rule that concludes that Coercible (Maybe τ) (Maybe σ), namely the lifting rule. As that rule has the assumption that Coercible $\tau \sigma$ holds, we can invert that rule and we can conclude Coercible $\tau \sigma$ from Coercible (Maybe τ) (Maybe σ). In other words, Maybe is injective with respect to coercibility. This is the *decomposition rule* (elided from Figure 3), and it can be used for any parameter of a non-**newtype** type constructor, e.g.:

- If Coercible (Maybe τ) (Maybe σ) then Coercible $\tau \sigma$.
- If Coercible $[\tau] [\sigma]$ then Coercible $\tau \sigma$.
- If Coercible (Either $\tau_1 \tau_2$) (Either $\sigma_1 \sigma_2$) then Coercible $\tau_1 \sigma_1$.
- If Coercible (Either $\tau_1 \tau_2$) (Either $\sigma_1 \sigma_2$) then Coercible $\tau_2 \sigma_2$.
- If Coercible $(\tau_1 \rightarrow \tau_2) (\sigma_1 \rightarrow \sigma_2)$ Coercible $\tau_1 \sigma_1$.
- If Coercible $(\tau_1 \rightarrow \tau_2) (\sigma_1 \rightarrow \sigma_2)$ Coercible $\tau_2 \sigma_2$.

¹ Or data family. See Chakravarty *et al.* (2005b) for more information.

The general rule follows:

- Suppose non-**newtype** T has parameters with roles **representational**, **phantom**, and **nominal**, respectively. If $\text{Coercible } (T \tau_1 \tau_2 \tau_3) (T \sigma_1 \sigma_2 \sigma_3)$, then $\text{Coercible } \tau_1 \sigma_1$ and $\tau_3 \sim \sigma_3$, where \sim is Haskell's notation for type equality.

The practical impact of adding the decomposition rule is small, as programs that require this rule are likely very rare. Nevertheless, the rule is sound and such programs might exist, so we include it. Furthermore, the formalism does need this decomposition in order to define a type-safe small-step semantics (cf. Appendix A.3).

Why do we have to exclude **newtypes** from the decomposition rule? Although they have lifting rules, Coercible constraints between them could also have been created using the unwrapping rule, so the argument by rule inversion above does not hold. In other words, **newtypes** are *not* injective with respect to coercibility. Indeed, if we assumed such a decomposition rule, we could derive invalid Coercible constraints. Consider the following code, where the programmer explicitly uses a role annotation (see Section 3.1) to set the role of the argument to **representational**:

```
newtype TaggedInt a = MkTI Int
type role TaggedInt representational
```

The explicit role annotation is fine, as explained in Section 4.3. Using the unwrapping rule, together with transitivity and symmetry, we can conclude that $\text{Coercible } (\text{TaggedInt Bool}) (\text{TaggedInt Char})$ holds. If we had a decomposition rule, we would now have $\text{Coercible Bool Char}$.

2.8 Supporting higher-order polymorphism

So far, we have only seen Coercible applied to types of kind $*$, but that is not sufficient to support all coercions that we might want. For example, consider a monad transformer such as

```
data MaybeT m a = MaybeT (m (Maybe a))
```

and a newtype that wraps another monad, e.g.

```
newtype MyIO a = MyIO (IO a)
```

It is reasonable to expect that $\text{Coercible } (\text{MaybeT MyIO a}) (\text{MaybeT IO a})$ can be derived. Using the lifting rule for MaybeT, this requires Coercible MyIO IO to hold. Therefore, for a **newtype** declaration as the one above, GHC will η -reduce the unwrapping rule to say Coercible IO MyIO instead of $\text{Coercible (IO a) (MyIO a)}$. Using symmetry, this allows us to solve $\text{Coercible } (\text{MaybeT MyIO a}) (\text{MaybeT IO a})$.

Of course, this η -reduction must not prevent us from solving, for example, $\text{Coercible } (\text{MyIO Int}) (\text{IO Int})$. Therefore, we have the *type application rule* that allows us to use Coercible relations between types of higher kinds such as $* \rightarrow *$:

- If $\text{Coercible } \tau \sigma$, where $\tau, \sigma :: \kappa_1 \rightarrow \kappa_2$, then $\text{Coercible } (\tau \tau_0) (\sigma \tau_0)$ for any τ_0 .

What about the very similar-looking rule “If Coercible $\tau \sigma$ then Coercible $(\tau_0 \tau)$ $(\tau_0 \sigma)$ ”, where τ_0 is a type variable (of kind $\kappa_1 \rightarrow \kappa_2$)? Such a lifting rule for type variables would be unsound. For example, the variable could be instantiated with a type constructor that has a nominal parameter, such as `EncText`, which would allow us to coerce (erroneously) between `EncText HTML` and `EncText String`.

Therefore, the parameters of type variables are always assumed to have nominal role, and no lifting rule is available. This inability to abstract over types while retaining information about their parameter’s roles has some practical consequences; see Section 8.

3 Roles, abstraction, and coherence

The purpose of the `HTML` type from the introduction is to prevent the confusion of unescaped strings and `HTML` fragments. However, because these types have the same representation, confusing them does not lead to unsoundness in the type system. Instead, programs that make this mistake do not preserve the user-defined abstraction of the `HTML` type.

While the previous section describes how the Coercible formation rules ensure that Coercible types share runtime representations, this section discusses two other properties that guide the design of this mechanism: *type abstraction* (Section 3.1) and *class coherence* (Section 3.2).

3.1 Preserving abstraction

Haskell programmers define *abstract types* by hiding the constructors of newtypes and datatypes. In this case, the creation of values of a type like `HTML` is controlled by a code in a single module, so programmers can establish invariants about those values. Because `coerce` can also construct values of type `HTML`, the unwrapping coercion associated with this newtype is available if and only if the newtype constructor `MkHTML` is in scope.

However, what about the interaction between the *lifting rule* and type abstraction? It turns out that, even when module authors have carefully hidden the constructors of a type, sometimes they want to make the lifting rule available for that type, but at other times they would like to restrict it.

To illustrate the former case, we would like to permit coercions between `IORef HTML` to `IORef String`, even though `IORef` is an abstract type. Similarly, consider a library for non-empty lists:

```
module NonEmptyListLib( NE, singleton, ... ) where
```

```
data NE a = MkNE [a]
singleton :: a → NE a
...etc...
```

The type must be exported abstractly; otherwise, the non-empty property can be broken by its users. Nevertheless lifting a coercion through `NE`, i.e. coercing `NE HTML` to `NE String`, does not break this invariant.

To illustrate the case where one would want to restrict the lifting rule, consider the data type `Map k v`. This type implements an efficient finite map from keys of type `k` to values of type `v` using an internal representation based on a balanced tree, something like this:

```
data Map k v = Leaf | Node k v (Map k v) (Map k v)
```

It would be disastrous if the user were allowed to coerce from `(Map Age v)` to `(Map Int v)`, because a valid tree with regard to the ordering of `Age` might be bogus when using the ordering of `Int`. Functions that manipulate Maps use an `Ord k` constraint and thus use the `Ord` instance for the type `k`; nothing in Haskell requires that instances `Ord Int` and `Ord Age` behave similarly.

To prevent coercing `(Map Age v)` to `(Map Int v)`, the programmer would explicitly give a *role annotation* that differs from the default annotation, namely:

```
type role Map nominal representational
```

As explained in Section 2.2, these roles produce the abstraction-preserving lifting rule

► If `Coercible a b` then `Coercible (Map k a) (Map k b)`

which allows the coercion from `Map k HTML` to `Map k String`.

Note that in the declaration of `Map` the parameters `k` and `v` are used in exactly the same way, so this distinction cannot be made by the compiler; it can only be specified by the programmer. However, the compiler ensures that programmer-specified role annotations cannot violate type safety: if the annotation specifies an unsafe role, the compiler will reject the program.

3.2 Preserving class coherence

Another property of Haskell, independent of type-safety, is the coherence of type classes. There should only ever be one class instance for a particular class and type. We call this desirable property *coherence*. Without careful design coercion could be used to create incoherence.

To demonstrate that, consider the type class `Show` and two of its instances:

```
class Show a where
```

```
  show :: a → String
```

```
instance Show String where
```

```
  show s = "\ " ++ s ++ "\ "
```

```
instance Show HTML where
```

```
  show (MkHTML s) = "<html>" ++ s ++ "</html>"
```

The following (non-Haskell98) data type reifies a `Show` instance as a value:

```
data HowToShow a where
```

```
  MkHTS :: Show a ⇒ HowToShow a
```

```
showH :: HowToShow a → a → String
```

```
showH MkHTS x = show x
```

Here `showH` pattern-matches on a `HowToShow` value, and uses the instance stored inside it to obtain the `show` method.

If we are not careful, we could break the coherence of the `Show` type class:

```
stringShow :: HowToShow String
stringShow = MkHTS
htmlShow :: HowToShow HTML
htmlShow = MkHTS
badShow :: HowToShow HTML
badShow = coerce stringShow
```

```
λ> putStrLn (showH stringShow "Hello")
"Hello"
λ> putStrLn (showH htmlShow (MkHTML "Hello"))
<html>Hello</html>
λ> putStrLn (showH badShow (MkHTML "Hello"))
"Hello"
```

In the last interaction we applied `show` to a value of type `HTML`, but the `Show` instance for `String` (coerced to `(Show HTML)`) was used. This example shows the problem that derives from the lack of coherence – we used `coerce` to construct a second instance of the `Show` class for the `HTML` type.

To avoid this, the parameters of a type class can only be assigned a *nominal* role.² Accordingly, the parameter of `HowToShow` is also assigned a nominal role, preventing the coercion between `(HowToShow HTML)` and `(HowToShow String)`.

4 Ensuring type safety: System FC with roles

Haskell is a large and complicated language. How do we know that the ideas sketched above in source language terms actually produce a sound type system? What, precisely, do roles mean, and when precisely are two types equal? In this section we answer these questions for GHC’s small, statically-typed intermediate language, `GHC Core`. Every Haskell program is translated into `Core`, and we can typecheck `Core` to reassure ourselves that the (large, complicated) front end accepts only good programs.

`Core` is an implementation of a calculus called `System FC`, itself an extension of the classical Girard/Reynolds `System F`. The version of `FC` that we develop in this paper derives from much prior work.³ However, for clarity we give a self-contained description of the system and do not assume familiarity with previous versions.

² A role annotation can be used to override this default, but the user must specify GHC’s `-XIncoherentInstances` extension to do so.

³ Several versions of `System FC` are described in published work. Some of these variants have had decorations to the `FC` name, such as `FC2` or `FC†`. We do not make these distinctions in the present work, referring instead to all of these systems – in fact, one evolving system – as “`FC`”.

Metavariables:			
x	term	a, b type	c coercion
C	axiom	D data type	N newtype
F	type family	K data constructor	
e	$::= \lambda c: \phi. e \mid e \gamma \mid e \triangleright \gamma \mid \dots$		terms
τ, σ	$::= a \mid \tau_1 \tau_2 \mid \forall a: \kappa. \tau \mid H \mid F(\bar{\tau})$		types
κ	$::= \star \mid \kappa_1 \rightarrow \kappa_2$		kinds
H	$::= (\rightarrow) \mid (\Rightarrow) \mid (\sim_{\rho}^{\kappa}) \mid T$		type constants
T	$::= D \mid N$		algebraic data types
ϕ	$::= \tau \sim_{\rho}^{\kappa} \sigma$		propositions
γ, η	$::=$		coercions
	$\mid \langle \tau \rangle \mid \langle \tau, \sigma \rangle_{\rho} \mid \mathbf{sym} \gamma \mid \gamma_1 \circ \gamma_2$		equivalence
	$\mid H(\bar{\gamma}) \mid F(\bar{\gamma}) \mid \gamma_1 \gamma_2 \mid \forall a: \kappa. \gamma$		congruence
	$\mid c \mid C(\bar{\tau})$		assumption
	$\mid \mathbf{nth}^i \gamma \mid \mathbf{left} \gamma \mid \mathbf{right} \gamma \mid \gamma @ \tau$		decomposition
	$\mid \mathbf{sub} \gamma$		sub-rolling
ρ	$::= \mathbf{N} \mid \mathbf{R} \mid \mathbf{P}$		roles
Γ	$::= \emptyset \mid \Gamma, a: \kappa \mid \Gamma, c: \phi \mid \Gamma, x: \tau$		typing contexts
Ω	$::= \emptyset \mid \Omega, a: \rho$		role contexts

Figure 4. An excerpt of the grammar of System FC

Figure 4 gives the syntax of System FC. The starting point is a conventional typed, polymorphic lambda calculus inspired by System F with algebraic datatypes. We therefore elide most of the syntax of terms e , giving the typing judgement for terms in Appendix A.2. Types τ are also conventional, except that we add (saturated) type-family applications $F(\bar{\tau})$, to reflect their addition to source Haskell (Chakravarty *et al.*, 2005a; Chakravarty *et al.*, 2005b).⁴ Types are classified by kinds κ as usual; the kinding judgement $\Gamma \vdash \tau : \kappa$ appears in Appendix A.2. This judgement is syntax directed: from the context Γ and type τ , we can determine the unique kind κ (if one exists). To avoid clutter we use only monomorphic kinds, but it is easy to add kind polymorphism along the lines of Yorgey *et al.* (2012), and our implementation does so.

FC is an *explicitly-typed* language. By using System F’s explicit type abstraction and application, an FC program can be typechecked by a simple, syntax-directed algorithm, despite the presence of impredicative polymorphism. Type inference is not required.

⁴ Type families must always appear saturated – even in System FC – because otherwise we would be unable to admit the **left** and **right** coercions. See also Section 4.2.7.

4.1 Roles and casts

FC’s distinctive feature is a type-safe cast ($e \triangleright \gamma$) (Figure 4), which uses a *coercion* γ to cast a term from one type to another. The explicit coercions and casts in System FC ensure that type checking remains simple and syntax-directed, despite the presence of GADTs and type families.

A coercion γ is a witness or proof of the equality of two types. Coercions are classified by the judgement

$$\Gamma \vdash \gamma : \tau \sim_{\rho}^{\kappa} \sigma$$

given in Figure 5, and pronounced “in type environment Γ the coercion γ witnesses that the types τ and σ both have kind κ , and are equal at role ρ ”.

The notion of being “equal at role ρ ” is the important feature of this paper; it is a development of earlier work, as Section 7 describes. There are precisely three roles (see Figure 4), written N, R, and P, with the following meaning:

Nominal equality, written \sim_N , is the equality that the source Haskell type checker reasons about. When a Haskell programmer says that two Haskell types are the “same”, we mean that the types are nominally equal. Thus, we can say that $\text{Int} \sim_N \text{Int}$ but **not** $\text{Int} \sim_N \text{Age}$. Type families introduce new nominal equalities. So, if we have **type instance** $F \text{ Int} = \text{Bool}$, then $F \text{ Int} \sim_N \text{Bool}$.

Representational equality, written \sim_R , holds between two types that share the same run-time representation. Because all types that are nominally equal also share the same representation, nominal equality is a subset of representational equality. Continuing the example from the introduction, $\text{HTML} \sim_R \text{String}$. A Coercible constraint in Haskell corresponds to a proposition of representational equality in FC.

Phantom equality, written \sim_P , holds between any two types, whatsoever. It may seem odd that we produce and consume proofs of this “equality”, but doing so keeps the system uniform and easier to reason about. The idea of phantom equality is new in this work, and it allows for zero-cost conversions among types with phantom parameters.

We can now give the typing rule for type-safe cast:

$$\frac{\Gamma \vdash e : \tau_1 \quad \Gamma \vdash \gamma : \tau_1 \sim_R \tau_2}{\Gamma \vdash e \triangleright \gamma : \tau_2} \quad \text{TM_CAST}$$

The coercion γ must be a proof of *representational* equality, as witnessed by the R subscript to the result of the coercion typing premise. This makes sense: we can treat an expression of one type τ_1 as an expression of some other type τ_2 when those types share a representation.

$$\begin{array}{c}
\frac{\Gamma \vdash \gamma : \tau \sim_{\mathbf{N}} \sigma}{\Gamma \vdash \mathbf{sub} \gamma : \tau \sim_{\mathbf{R}} \sigma} \text{ CO_SUB} \quad \frac{\Gamma \vdash \gamma : \sigma \sim_{\rho} \tau}{\Gamma \vdash \mathbf{sym} \gamma : \tau \sim_{\rho} \sigma} \text{ CO_SYM} \\
\frac{\Gamma \vdash \gamma_1 : \tau_1 \sim_{\rho} \tau_2 \quad \Gamma \vdash \gamma_2 : \tau_2 \sim_{\rho} \tau_3}{\Gamma \vdash \gamma_1 \mathbin{\&}; \gamma_2 : \tau_1 \sim_{\rho} \tau_3} \text{ CO_TRANS} \\
\frac{\Gamma \vdash \tau : \kappa}{\Gamma \vdash \langle \tau \rangle : \tau \sim_{\mathbf{N}} \tau} \text{ CO_REFL} \quad \frac{\Gamma \vdash \tau : \kappa \quad \Gamma \vdash \sigma : \kappa}{\Gamma \vdash \langle \tau, \sigma \rangle_{\rho} : \tau \sim_{\rho} \sigma} \text{ CO_PHANTOM} \\
\frac{C : [\bar{a} : \bar{\kappa}]. \sigma_1 \sim_{\rho} \sigma_2}{\vdash \Gamma \quad \Gamma \vdash \tau : \kappa} \text{ CO_AXIOM} \quad \frac{\vdash \Gamma \quad c : \tau \sim_{\rho} \sigma \in \Gamma}{\Gamma \vdash c : \tau \sim_{\rho} \sigma} \text{ CO_VAR} \\
\frac{\Gamma \vdash \gamma : \tau \sim_{\rho} \sigma \quad \bar{\rho} \text{ is a prefix of roles}(H)}{\Gamma \vdash H \bar{\tau} : \kappa \quad \Gamma \vdash H \bar{\sigma} : \kappa} \text{ CO_TYCONAPP} \\
\frac{\Gamma \vdash \gamma_1 : \tau_1 \sim_{\rho} \sigma_1 \quad \Gamma \vdash \gamma_2 : \tau_2 \sim_{\mathbf{N}} \sigma_2}{\Gamma \vdash \tau_1 \tau_2 : \kappa \quad \Gamma \vdash \sigma_1 \sigma_2 : \kappa} \text{ CO_APP} \\
\frac{\Gamma \vdash \gamma : \tau \sim_{\mathbf{N}} \sigma \quad \Gamma \vdash F(\bar{\tau}) : \kappa \quad \Gamma \vdash F(\bar{\sigma}) : \kappa}{\Gamma \vdash F(\bar{\gamma}) : F(\bar{\tau}) \sim_{\mathbf{N}} F(\bar{\sigma})} \text{ CO_TYFAM} \\
\frac{\Gamma, a : \kappa \vdash \gamma : \tau \sim_{\rho} \sigma}{\Gamma \vdash \forall a : \kappa. \gamma : \forall a : \kappa. \tau \sim_{\rho} \forall a : \kappa. \sigma} \text{ CO_FORALL} \\
\frac{\Gamma \vdash \gamma : H \bar{\tau} \sim_{\mathbf{R}} H \bar{\sigma} \quad \bar{\rho} \text{ is a prefix of roles}(H) \quad H \text{ is not a newtype}}{\Gamma \vdash \mathbf{nth}^i \gamma : \tau_i \sim_{\rho_i} \sigma_i} \text{ CO_NTH} \\
\frac{\Gamma \vdash \gamma : \tau_1 \tau_2 \sim_{\mathbf{N}} \sigma_1 \sigma_2 \quad \Gamma \vdash \tau_1 : \kappa \quad \Gamma \vdash \sigma_1 : \kappa}{\Gamma \vdash \mathbf{left} \gamma : \tau_1 \sim_{\mathbf{N}} \sigma_1} \text{ CO_LEFT} \quad \frac{\Gamma \vdash \gamma : \tau_1 \tau_2 \sim_{\mathbf{N}} \sigma_1 \sigma_2 \quad \Gamma \vdash \tau_2 : \kappa \quad \Gamma \vdash \sigma_2 : \kappa}{\Gamma \vdash \mathbf{right} \gamma : \tau_2 \sim_{\mathbf{N}} \sigma_2} \text{ CO_RIGHT} \\
\frac{\Gamma \vdash \gamma : \forall a : \kappa. \tau_1 \sim_{\rho} \forall a : \kappa. \sigma_1 \quad \Gamma \vdash \tau : \kappa}{\Gamma \vdash \gamma @ \tau : \tau_1[\tau/a] \sim_{\rho} \sigma_1[\tau/a]} \text{ CO_INST}
\end{array}$$

Figure 5. $\Gamma \vdash \gamma : \phi$: Formation rules for coercions

4.2 Coercions

Coercions (Figure 4) and their typing rules (Figure 5) are the heart of System FC. The basic typing judgement for coercions is $\Gamma \vdash \gamma : \tau \sim_{\rho}^{\kappa} \sigma$. This judgement is also syntax directed: when this judgement holds we can determine the unique proposition $\tau \sim_{\rho}^{\kappa} \sigma$ that is justified by a particular coercion γ in a given context Γ . Furthermore, in this case τ and σ must be well formed and have the same kind κ . We often omit the kind annotation in our presentation when it is not important.

We can understand the typing rules in Figure 5, by thinking about the equalities that they define:

4.2.1 Nominal equality implies representational equality

If we have a proof that two types are nominally equal, then they are in particular also representationally equal. This intuition is expressed by the **sub** operator, and the rule CO_SUB. With this, the above type-safe cast using TM_CAST can also make use of nominal equalities.

4.2.2 Phantom equality relates all types

The coercion form $\langle \tau, \sigma \rangle_P$ (shown in rule CO_PHANTOM) proves that any two types τ and σ are equal at role P.

4.2.3 Equality is an equivalence relation

Equality is an equivalence relation at all three roles. Symmetry (rule CO_SYM) and transitivity (CO_TRANS) work for any role ρ . Reflexivity is modelled directly only for nominal equality, by way of CO_REFL. It also holds at the other roles: we can derive representational reflexivity using **sub**, and phantom equality trivially includes reflexivity through rule CO_PHANTOM.

4.2.4 Axioms for equality

Each newtype declaration and type-family instance gives rise to an *axiom*; newtypes give rise to representational axioms, and type-family instances give rise to nominal axioms.⁵ For example, the declarations

```
newtype HTML = MkHTML String
newtype EitherInt a = MkEI (Either a Int)
type family F [a] = Maybe a
```

produce the axioms

$$\begin{aligned} C_1 &: \text{HTML} \sim_R \text{String} \\ C_2 &: [a:\star]. \text{EitherInt } a \sim_N \text{Either } a \text{ Int} \\ C_3 &: [a:\star]. F([a]) \sim_N \text{Maybe } a \end{aligned}$$

Axiom C_1 states that HTML is *representationally* equal to String, just as Axiom C_2 states that $\text{EitherInt } \sigma$ is *representationally* equal to $\text{Either } \tau \text{ Int}$ for any type σ : they are distinct types, but share a common representation. Axiom C_3 states that $F([\sigma])$ is *nominally* equal to $\text{Maybe } \sigma$: the two are considered to be the same type by the type checker).

In C_2 and C_3 , the notation “[$a:\star$].” binds a in the types being equated. Uses of these axioms are governed by the rule CO_AXIOM. Axioms must always appear fully applied, and we assume that they live in a global context, separate from the local context Γ .

⁵ For simplicity, we restrict ourselves to *open* type families. Closed type families (Eisenberg et al., 2014) could also be accommodated.


```

newtype HTML = MkHTML String

type family F a
type instance F String = Int
type instance F HTML = Bool
— F’s parameter has a nominal role

data T a = MkT (F a)
type role T nominal

```

Figure 6. Congruence and roles example code

4.2.5 Equality can be abstracted

Just as one can abstract over types and values in System F, one can also abstract over equality proofs in FC. To this end, FC terms (Figure 4) include coercion abstraction $\lambda c:\phi.e$ and application $e\gamma$. These are the introduction and elimination forms for the coercion abstraction arrow (\Rightarrow), just as ordinary value abstraction and application are the introduction and elimination forms for ordinary arrow (\rightarrow) (see Appendix A.2).

A coercion abstraction binds a coercion variable $c:\phi$. These variables can occur only in coercions; see rule CO_VAR. Coercion variables can also be bound in the patterns of a **case** expression, which supports generalised algebraic data types (GADTs).

4.2.6 Equality is congruent

Congruence of type application Before diving into the rules themselves, it is helpful to consider some examples of how we want congruence and roles to interact. Let’s consider the definitions in Figure 6. The role annotation here is superfluous, as nominal is the only legal role for T. With these definitions in hand, what equalities should be derivable? (Recall the intuitive meanings of the different roles in Section 4.1.)

1. Should $\text{Maybe HTML} \sim_R \text{Maybe String}$ hold?

Yes, it should. The type parameter to Maybe has a representational role, so it makes sense that two Maybes built out of representationally equal types should be representationally equal.
2. Should $\text{Maybe HTML} \sim_N \text{Maybe String}$ hold?

Certainly not. These two types are entirely distinct to Haskell programmers and its type checker.
3. Should $T \text{ HTML} \sim_R T \text{ String}$ hold?

Certainly not. We can see, by unfolding the definition for T, that the representations of the two types are different.
4. Should $a \text{ HTML} \sim_R a \text{ String}$ hold, for a type variable a ?

It depends on the instantiation of a ! If a becomes Maybe, then “yes”; if a becomes T, then “no”. Since we may be abstracting over a , we do not know

which of the two will happen, so we take the conservative stance and say that $a \text{ HTML} \sim_{\mathbb{R}} a \text{ String}$ does *not* hold.

This last point is critical. The alternative is to express a 's argument roles in its kind, but that leads to a more complicated system; see related work in Section 7. A distinguishing feature of this paper is the simplification we obtain by attributing roles only to the arguments to type constants (H , in the grammar), and not to abstracted type variables. We lose a little expressiveness; see Sections 7.1 and 8.1 for discussion and examples.

To support both (1) and (4) requires two coercion forms and corresponding rules:

- The coercion form $H(\bar{\gamma})$ has an explicit type constant at its head. This form always proves a representational equality, and it requires input coercions of the roles designated by the roles of H 's parameters (rule CO_TYCONAPP). The *roles* function gives the list of roles assigned to H 's parameters, as explained in Section 2.2. We allow $\bar{\rho}$ to be a prefix of *roles*(H) to accommodate partially-applied type constants. The $H(\bar{\gamma})$ form allows coercions like in case (1) but not case (4).
- The coercion form $\gamma_1 \gamma_2$ does not have an explicit type constant, so we must use the conservative treatment of roles discussed above. Rule CO_APP therefore requires γ_2 to be a nominal coercion, though the role of γ_1 carries through to the application $\gamma_1 \gamma_2$. This form addresses case (4).

What if we wish to prove a nominal equality such as $\text{Maybe}(\text{FString}) \sim_{\mathbb{N}} \text{MaybeInt}$? We can't use the $H(\bar{\gamma})$ form, which proves only representational equality, but we can use the $\gamma_1 \gamma_2$ form, with $\langle \text{Maybe} \rangle$ for γ_1 .

Congruence of type family application Rule CO_TYFAM proves the equality of two type-family applications. It requires nominal coercions among all the arguments because type families can inspect their (type) arguments and branch on them. It would be unsound to derive an equality between FString and FHTML .

Congruence of polymorphic types The rule CO_FORALL works for any role ρ ; polymorphism and roles do not interact.

4.2.7 Equality can be decomposed

If we have a proof of $\text{Maybe} \sigma \sim_{\rho} \text{Maybe} \tau$, with type applications on both sides, should we be able to get a proof of $\sigma \sim_{\rho} \tau$, by decomposing the equality? Yes, in this case, but we must be careful here as well.

The formation rule CO_TYCONAPP has an (almost) inverse in the rule CO_NTH, and the formation rule CO_APP has an (almost) inverse in the rules CO_LEFT and CO_RIGHT; these rule can decompose some type applications. For various reasons, though, these must not be complete inverses:

- The rule CO_NTH is not allowed to decompose equalities among newtypes. Why? Because **nth** witnesses injectivity and newtypes are not necessarily injective with respect to representational equality. (Like all datatypes in Haskell, newtypes are injective with respect to nominal equality.) For example, consider these definitions:

```
data Phant a = MkPhant
type role Phant phantom
newtype App a b = MkApp (a b)
type role App representational nominal
```

So we have, $roles(App) = R, N$. Yet, we can see the following chain of equalities:

$$App\ Phant\ Int \sim_R Phant\ Int \sim_R Phant\ Bool \sim_R App\ Phant\ Bool$$

By transitivity, we can derive a coercion γ witnessing

$$App\ Phant\ Int \sim_R App\ Phant\ Bool$$

If we could use **nth**² on γ , we would get $Int \sim_N Bool$: disaster! We eliminate this possibility by preventing **nth** on newtypes.

- The rules CO_LEFT and CO_RIGHT require and produce only nominal coercions, and are not allowed to operate on representational coercions. Consider this newtype to see why this must be so:

```
newtype EitherInt a = MkEI (Either a Int)
```

This definition yields an axiom showing that, for all a , $EitherInt\ a \sim_R (Either\ a\ Int)$. Suppose we could apply **left** and **right** to coercions formed from this axiom. Using **left** would get us a proof of $EitherInt \sim_R (Either\ a)$, which could then be used to show, say, $(Either\ Char) \sim_R (Either\ Bool)$ and then (using **nth**) $Char \sim_N Bool$. Using **right** would get us a proof of $a \sim_R Int$, for *any* a . These are both clearly disastrous. So, we forbid using these coercion formers on representational coercions.⁶

- The rule CO_NTH is restricted to decompose only representational coercions, and never nominal ones. This is not because it would be unsound otherwise, but as nominal coercions can be decomposed via **left** and **right**, there is simply no need for using **nth** on nominal coercions, and as a design decision, we left them out.

Thankfully, polymorphism and roles play well together, and the CO_INST rule (inverse to CO_FORALL) shows quite straightforwardly that, if two polytypes are equal, then so are the instantiated types.

⁶ Although the forms **left** and **right** were originally part of FC, for simplicity they were omitted in previous papers (Weirich *et al.*, 2011) and in the implementation (GHC 7.2-7.6). However, Haskell users (e.g. Trac #7205) reported that some programs no longer type checked after this change, so these forms were re-introduced in GHC 7.8.

$$\boxed{\bar{\rho} \models H} \quad \text{“}\bar{\rho} \text{ are appropriate roles for } H\text{.”}$$

$$\frac{
\begin{array}{l}
\forall \bar{a}, \bar{b}, \bar{\sigma} \text{ s.t. } K : \forall \bar{a} : \bar{\kappa}. \forall \bar{b} : \bar{\kappa}'. \bar{\phi} \Rightarrow \bar{\sigma} \rightarrow D \bar{a}, \\
\forall \tau \text{ s.t. } \tau \in \bar{\sigma} \vee \tau \in \bar{\phi}, \\
\bar{a} : \bar{\rho}, \bar{b} : \bar{N} \vdash \tau : R
\end{array}
}{\bar{\rho} \models D} \quad \text{ROLES_DATA}$$

$$\frac{C : [\bar{a} : \bar{\kappa}]. N \bar{a} \sim_R \sigma \quad \bar{a} : \bar{\rho} \vdash \sigma : R}{\bar{\rho} \models N} \quad \text{ROLES_NEWTYP}$$

$$\overline{R, R} \models (\rightarrow) \quad \overline{R, R} \models (\Rightarrow) \quad \overline{\rho, \rho} \models (\sim \rho)$$

$$\boxed{\Omega \vdash \tau : \rho} \quad \text{“Assuming } \Omega, \tau \text{ can be used at role } \rho\text{.”}$$

$$\frac{a : \rho' \in \Omega \quad \rho' \leq \rho}{\Omega \vdash a : \rho} \quad \text{RTY_VAR}$$

$$\frac{\bar{\rho} \text{ is a prefix of } \text{roles}(H) \quad \overline{\Omega \vdash \tau : \rho}}{\Omega \vdash H \bar{\tau} : R} \quad \text{RTY_TYCONAPP}$$

$$\frac{}{\Omega \vdash H : N} \quad \text{RTY_TYCON}$$

$$\frac{\Omega \vdash \tau : \rho \quad \Omega \vdash \sigma : N}{\Omega \vdash \tau \sigma : \rho} \quad \text{RTY_APP}$$

$$\frac{\Omega, a : N \vdash \tau : \rho}{\Omega \vdash \forall a : \kappa. \tau : \rho} \quad \text{RTY_FORALL}$$

$$\frac{\overline{\Omega \vdash \tau : N}}{\Omega \vdash F(\bar{\tau}) : \rho} \quad \text{RTY_TYFAM}$$

$$\frac{}{\Omega \vdash \tau : P} \quad \text{RTY_PHANTOM}$$

$$\boxed{\rho_1 \leq \rho_2} \quad \text{“}\rho_1 \text{ is a sub-role of } \rho_2\text{.”}$$

$$\overline{N} \leq \overline{\rho} \quad \overline{\rho} \leq \overline{P} \quad \overline{\rho} \leq \overline{\rho}$$

Figure 7. Rules asserting a correct assignment of roles to data types

There is no decomposition form for type family applications: knowing that $F(\bar{\tau})$ is equal to $F(\bar{\sigma})$ tells us nothing whatsoever about the relationship between $\bar{\tau}$ and $\bar{\sigma}$.

4.3 Role attribution for type constants

In System FC we assume an unwritten global environment of top-level constants: data types, type families, axioms, and so on. For a data type H , for example, this

environment gives the kind of H , the types of H 's data constructors, and the roles of H 's parameters. Abstract data types, imported from other modules, are data types that include no data constructors in the global environment. Clearly this global environment must be internally consistent. For example, a data constructor K must return a value of type $D \bar{\tau}$ where D is a data type; K 's type must be well-kinded, and that kind must be consistent with D 's kind.

All of this is standard except for roles. It is essential that the roles of D 's parameters, $roles(D)$, are consistent with D 's definition. For example, it would be wrong for the global environment to claim that `Maybe`'s parameter is phantom because then we could use `CO_PHANTOM` and `CO_TYCONAPP` to prove the obviously wrong coercion `MaybeInt ~R MaybeBool`.

We use the judgement $\bar{\rho} \models H$, to mean “ $\bar{\rho}$ are suitable roles for the parameters of H ”, and in our proof of type safety, we assume that $roles(H) \models H$ for all H . The rules for this judgement and two auxiliary judgements appear in Figure 7. Note that this judgement defines a *relation* between roles and data types. Our role inference algorithm (Section 4.6) determines the most permissible roles for this relation, but often other, less permissive roles, such as those specified by role annotations, are also included by this relation.

Start with `ROLES_NEWTYPE`. Recall that a newtype declaration for N gives rise to an axiom $C : [\bar{a}:\bar{\kappa}]. N \bar{a} \sim_R \sigma$. The rule says that roles $\bar{\rho}$ are acceptable for N if each parameter a_i is used in σ in a way consistent with ρ_i , expressed using the auxiliary judgement $\bar{a}:\bar{\rho} \vdash \sigma : R$.

The key auxiliary judgement $\Omega \vdash \tau : \rho$ checks that the type variables in τ are used in a way consistent with their roles specified in Ω , when considered at role ρ . More precisely, the main purpose of the judgement is to guarantee that if $\Omega \vdash \tau : \rho$ holds, then for every $a:\rho' \in \Omega$ and $\sigma_1 \sim_{\rho'} \sigma_2$ we obtain $\tau[\sigma_1/a] \sim_{\rho} \tau[\sigma_2/a]$. This is a consequence of the lifting lemma (Lemma 33). Unlike in many typing judgements, the role ρ (as well as Ω) is an *input* to this judgement, not an output. With this in mind, the rules for the auxiliary judgement are straightforward. For example, `RTY_TYFAM` says that the argument types of a type family application are used at nominal role. The variable rule, `RTY_VAR`, allows a variable to be assigned a more restrictive role (via the sub-role judgement) than required, which is needed both for multiple occurrences of the same variable, and to account for role signatures. Note that rules `RTY_TYCONAPP` and `RTY_APP` overlap – this judgement is not syntax-directed.

Returning to our original judgement $\bar{\rho} \models H$, `ROLES_DATA` deals with algebraic data types D , by checking roles in each of its data constructors K . The type of a constructor is parameterised by universal type variables \bar{a} , existential type variables \bar{b} , coercions with types $\bar{\phi}$ – which are used when encoding GADTs – and term-level arguments with types $\bar{\sigma}$. For each constructor, we must examine each proposition ϕ and each term-level argument type σ , checking to make sure that each is used at a representational role. Why check for a representational role specifically? Because $roles$ is used in `CO_TYCONAPP`, which produces a representational coercion. In other words, we must make sure that each term-level argument appears at a rep-

representational role within the type of each constructor K for CO_TYCONAPP to be sound.

The function type constructors (\rightarrow) and (\Rightarrow) have representational roles: functions care about representational equality but never branch on the nominal identity of a type. (For example, functions always treat `HTML` and `String` identically.)

Finally we see that the roles of the arguments to an equality proposition match the role of the proposition:

- It cannot have a less strict role, as otherwise we could use $\text{Int} \sim_{\text{P}} \text{Bool}$ to change $\text{Int} \sim_{\text{N}} \text{Int}$ into $\text{Int} \sim_{\text{N}} \text{Bool}$, which would not be sound.
- On the other hand, there is no point in using a stricter role: The ability to use $\gamma_1 : \tau_1 \sim_{\rho} \tau_2$ to change $\gamma_2 : \tau_1 \sim_{\rho} \tau_3$ into $\tau_2 \sim_{\rho} \tau_3$ is already given by the coercion expression $\mathbf{sym} \gamma_1 \circ \gamma_2$.

These definitions lead to a powerful theorem:

Theorem (Roles assignment narrowing). *If $\bar{\rho} \models H$, where H is a data type or newtype, and $\bar{\rho}'$ is such that $\rho'_i \leq \rho_i$ (for $\rho_i \in \bar{\rho}$ and $\rho'_i \in \bar{\rho}'$), then $\bar{\rho}' \models H$.*

Proof. Straightforward induction on $\Omega \vdash \tau : \rho$. □

This theorem states that, given a sound role assignment for H , any more restrictive role assignment is also sound. This property of our system here is one of its distinguishing characteristics from our prior work on roles – see the end of Section 8.1.1 for discussion.

4.4 Progress and preservation

The preceding discussion gave several non-obvious examples where admitting *too many* coercions would lead to unsoundness. However, we must have *enough* coercions to allow us to make progress when evaluating a program. (For example, the **nth** decomposition coercion is necessary for the `S_KPUSH` rule of the operational semantics, shown in Appendix A.3.) Happily, we can be confident that we have enough coercions, but not too many, because we prove the usual progress and preservation theorems for System FC.

The full proof of type safety appears in the appendix; it exhibits no new proof techniques. The structure of the proofs follows previous work, such as Weirich *et al.* (2011) or Yorgey *et al.* (2012). Despite following previous work, we have made one structural difference: our operational semantics allows evaluation under type abstractions $\Lambda a:\kappa.e$. This choice echoes the behaviour in GHC and prevents type abstractions from interfering with evaluation. It does not complicate the theory.

A key step in the proof of progress is to prove *consistency*; that is, that no coercion can exist between, say, `Int` and `Bool`. This is done by defining a non-deterministic, role-directed rewrite relation on types and showing that the rewrite system is confluent⁷ and preserves type constants (other than newtypes) appearing in the

⁷ As in prior work (Eisenberg *et al.*, 2014), we ensure that the rewrite relation is confluent by restricting type families to have only linear patterns. If non-linear patterns were allowed

heads of types. We then prove that, if a coercion exists between two types τ_1 and τ_2 , these two types both rewrite to a type σ . We conclude then that τ_1 and τ_2 , if headed by a non-newtype type constant, must be headed by the same such constant.

4.5 Type erasure

We claim that coercions (both our new representational coercions and the older nominal ones) are “zero-cost”. By this, we mean that the handling of coercions and the casts do not interfere in any way with runtime evaluation. In order to prove this claim, we define an erased language which removes all types and coercions.

4.5.1 Coercion abstractions

The one subtlety in this claim is that we cannot properly erase coercion abstractions. For example, the following closed expression is well-typed:

$$\lambda c:\text{Bool} \sim_{\text{R}} \text{Int}.3 + (\text{True} \triangleright c)$$

Imagine erasing the coercion abstraction and cast. We would be left with $3 + \text{True}$, which is certainly not a closed expression we wish to evaluate. Accordingly, we retain coercion abstractions even when erasing coercions. We thus must also retain coercion application. Formally, type erasure (denoted $|e|$) includes the following two equations (among other equations that erase types and casts):

$$|\lambda c:\phi.e| = \lambda \bullet. |e| \quad |e \gamma| = |e| \bullet$$

Beyond the forms above, the erased language is just the untyped λ -calculus with data constructors and **case** expressions.

Naturally, we define $\lambda \bullet. |e|$ as a value in our erased language. It therefore seems conceivable that evaluation could get hung up on a coercion abstraction. (That is, a Haskell program could evaluate to a coercion abstraction, which is a normal form.) This is silly, though, because a Haskell programmer does not think in terms of coercion abstractions and would not expect this compiler-generated form to interrupt the execution of a program.

The way we can claim a type erasure property is that *type inference never produces a coercion abstraction*. Backing up this claim is beyond the scope of this article, but the elaboration from Haskell to System FC embodied in the type-checking never needs coercion abstractions. All coercion variables are bound in **case** matches only.

Why have coercion abstractions, then? They are useful during program transformations, when we might want to combine two different branches of a **case**

in type families (that is, with a repeated variable on the left-hand side), combined with non-termination, our rewrite system would not be confluent. Losing confluence does not necessarily threaten consistency – it just threatens the particular proof technique that we use. However, a more powerful proof appears to be an open problem in the term rewriting community. Specifically, a positive answer to open problem #79 of the Rewriting Techniques and Applications (RTA) conference would lead to a proof of consistency; see <http://www.win.tue.nl/rtaloop/problems/79.html>.

expression, for example. It is a soundness property of these transformations that they do not affect the final value of an expression; thus we can be sure that no Haskell program is ever elaborated into an optimised System FC program that evaluates to a coercion abstraction.

4.5.2 Statement of type erasure

We prove in Appendix H that there is a bisimulation between reduction in the original and the erased language. More precisely, we prove that

- whenever e reduces to e' , then either $|e|$ reduces to $|e'|$ or $|e| = |e'|$, and
- whenever $|e|$ reduces, e does also.

The first claim says that the erased λ -calculus faithfully implements System FC. The second says that type abstraction, in particular, does not hold up evaluation. This is because our version of FC reduces under type abstractions. Please see the appendix for the details.

4.6 Role inference

We have assumed throughout this discussion a global context where we can look up roles via $roles(H)$ and that these roles are appropriate, i.e. $roles(H) \models H$ for all H . We give here the algorithm that populates the environment $roles(H)$:

- Primitive type constructors like (\rightarrow) and (\sim_ρ^k) have predefined roles for their parameters (Figure 7).
- Type families (Section 2.5) have nominal roles for all parameters.
- The roles of **class**, **data** type, or **newtype** parameters are determined by a role inference algorithm, which we describe next.

The role inference algorithm is straightforward. At a high level, it starts with the role information of the built-in constants (\rightarrow) , (\Rightarrow) , and (\sim_ρ) , and propagates roles until it finds a fixpoint. In the description of the algorithm below, we assume a mutable environment; $roles(H)$ pulls a list of roles from this environment. Only after the algorithm is complete will $roles(H) \models H$ hold.

1. Populate $roles(T)$ (for all T) with user-supplied annotations; omitted role annotations default to phantom for **data** and **newtype** and to nominal for **class**. Other than this default, classes are treated identically to datatypes, as they are implemented in FC via datatypes representing dictionaries (Sulzmann et al., 2007; Hall et al., 1996). (See Section 6.4 for discussion about this choice of default.)
2. For every data type D , every constructor for that data type K , for every σ in a proposition of that constructor, run $walk(D, \sigma)$, and for every term-level argument type σ to that constructor run $walk(D, \sigma)$.
3. For every newtype N with representation type σ , run $walk(N, \sigma)$.
4. If the role of any parameter to any type constant changed in the previous steps, go to step 2.

5. For every T , check $\text{roles}(T)$ against a user-supplied annotation, if any. If these disagree, reject the program. Otherwise, $\text{roles}(T) \models T$ holds.

The procedure $\text{walk}(T, \sigma)$ is defined as follows, matching from top to bottom:

$$\begin{aligned} \text{walk}(T, a) &:= \text{mark the } a \text{ parameter to } T \text{ as } R, \text{ when } a \text{ is unmarked.} \\ \text{walk}(T, H \bar{\tau}) &:= \text{let } \bar{\rho} = \text{roles}(H); \\ &\quad \text{for every } i, 0 < i \leq \text{length}(\bar{\tau}): \\ &\quad \quad \text{if } \rho_i = N, \text{ then} \\ &\quad \quad \quad \text{mark all variables free in } \tau_i \text{ as } N; \\ &\quad \quad \quad \text{else if } \rho_i = R, \text{ then } \text{walk}(T, \tau_i). \\ \text{walk}(T, \tau_1 \tau_2) &:= \text{walk}(T, \tau_1); \\ &\quad \quad \text{mark all variables free in } \tau_2 \text{ as } N. \\ \text{walk}(T, F(\bar{\tau})) &:= \text{mark all variables free in the } \bar{\tau} \text{ as } N. \\ \text{walk}(T, \forall b:\kappa. \tau) &:= \text{walk}(T, \tau). \end{aligned}$$

When marking variables, we ignore those that are not parameters to the data type T in question or have previously been marked as N . The first case deals with existential and local (\forall -bound) type variables and the second with the case where a variable is used both in a nominal and in a representational context.

Theorem. *The role inference algorithm always terminates.*

Theorem (Role inference is sound). *After running the role inference algorithm, $\text{roles}(H) \models H$ will hold for all H .*

Theorem (Role inference is optimal). *Suppose H has no role annotation. After running the role inference algorithm, any loosening of the roles assigned to H (a change from ρ to ρ' , where $\rho \leq \rho'$ and $\rho \neq \rho'$) would violate $\text{roles}(H) \models H$.*

Theorem (Role annotations only tighten roles). *Suppose a role annotation assigns roles $\bar{\rho}$ to H . If roles $\bar{\rho}'$ were inferred for a definition H' identical to H but missing H 's role annotation, then $\bar{\rho} \leq \bar{\rho}'$.*

Theorem (Principal role assignments). *For a given set of type constants \overline{H} , there is at most one choice of role assignments $\overline{\text{roles}(H)}$ that is optimal and such that $\overline{\text{roles}(H)} \models \overline{H}$.*

Arguments supporting these claims appear in Appendix G.

5 Type inference with Coercible constraints

Section 2 describes a programmer-level view of when types are Coercible; this section describes the portion of GHC's type inference algorithm that solves these constraints. This algorithm also produces the coercion evidence as described in Section 4, but we elide the details of evidence creation as this process is straightforward.

Type inference in GHC is accomplished via the OUTSIDEIN(X) algorithm, as described by Vytiniotis *et al.* (2011). This algorithm is a constraint-based type inference algorithm (Pottier & Rémy, 2005), that first generates a set of constraints

L		metavariable for classes
ξ	$::= a \mid \xi_1 \xi_2 \mid H \mid \forall a:\kappa.\xi$	function-free types
Q	$::= \epsilon \mid Q_1 \wedge Q_2 \mid L \bar{\tau} \mid \tau_1 \sim_N \tau_2 \mid \tau_1 \sim_R \tau_2$	constraints
Q	$::= Q \mid Q_1 \wedge Q_2$	top-level axiom schemes
	$\mid \forall \bar{a}:\bar{\kappa}.Q \Rightarrow L \bar{\tau}$	(constrained) class instance
	$\mid \forall \bar{a}:\bar{\kappa}.F(\bar{\tau}) \sim_N \sigma$	type family instance
	$\mid \forall \bar{a}:\bar{\kappa}.N \bar{a} \sim_R \sigma$	newtype axiom
ℓ	$::= g \mid w$	constraint flavours

Figure 8. Grammar for our constraint system

$Q \wedge Q \Vdash Q$	reflexivity
$Q \wedge Q_1 \Vdash Q_2$ and $Q \wedge Q_2 \Vdash Q_3$ implies $Q \wedge Q_1 \Vdash Q_3$	transitivity
$Q \Vdash Q_2$ implies $\theta(Q) \Vdash \theta(Q_2)$	substitutivity
$Q \Vdash \tau \sim_N \tau$	nominal eq. reflexivity
$Q \Vdash \tau_1 \sim_N \tau_2$ implies $Q \Vdash \tau_2 \sim_N \tau_1$	nominal eq. symmetry
$Q \Vdash \tau_1 \sim_N \tau_2$ and $Q \Vdash \tau_2 \sim_N \tau_3$ implies $Q \Vdash \tau_1 \sim_N \tau_3$	nominal eq. transitivity
$Q \Vdash Q_1$ and $Q \Vdash Q_2$ implies $Q \Vdash Q_1 \wedge Q_2$	conjunctions
$Q \Vdash \tau_1 \sim_N \tau_2$ implies $Q \Vdash \tau[\tau_1/a] \sim_N \tau[\tau_2/a]$	nominal eq. congruence

Figure 9. Requirements of the entailment relation $Q \Vdash Q$, adapted from Figure 3 of Vytiniotis *et al.* (2011).

during a pass over the Haskell source code and then solves these constraints separately. `OUTSIDEIN(X)` is parameterised by a constraint language and associated constraint solver; the X in `OUTSIDEIN(X)`. Our work fits into this framework by introducing a new Coercible `t1 t2` constraint and extending the constraint solver to handle this constraint.

5.1 A constraint system with representational equality

The grammar for our instantiation of X appears in Figure 8. A constraint Q can be empty (trivially satisfied), a conjunction of constraints, a class constraint $L \bar{\tau}$, or an equality constraint. A nominal equality constraint $\tau_1 \sim_N \tau_2$ is the standard type equality constraint already present in `OUTSIDEIN(X)`; a representational one $\tau_1 \sim_R \tau_2$ is the encoding of Coercible `\tau1 \tau2`. (Phantom equality constraints $\tau_1 \sim_P \tau_2$ are unnecessary.)

The constraint system X defines an entailment relation $Q \Vdash Q$, a judgement that holds whenever the assumptions Q imply the constraint Q . Note that the grammar for Q includes a conjunction of both regular constraints Q as well as top-level axioms. In our case, these axioms take one of three forms as shown in Figure 8: a class instance, a type family instance, or a newtype axiom. The `OUTSIDEIN(X)`

framework expects the entailment relation to uphold the properties listed in Figure 9.

In our case, the entailment relation essentially duplicates Figure 5, leaving out the form of the coercions themselves. Added onto those rules are rules for type classes, which do not concern us here. It can easily be shown that this entailment relation satisfies the properties of Figure 9. In particular, note that the substitutivity property of entailment is directly implied by a standard substitution lemma over coercions.

5.2 An overview of OUTSIDEIN(X)

We start with a brief overview of the OUTSIDEIN(X) algorithm, somewhat simplified from its original presentation.⁸ Our goal is not to provide a complete explanation of OUTSIDEIN(X), but to provide enough context to explain the modifications required by the new Coercible constraint. Due to the complexities they add to the algorithm, polytypes (headed by \forall) are excluded from this presentation; their complexity is orthogonal to roles.

OUTSIDEIN(X) uses a judgement $\Gamma \vdash e : \tau \rightsquigarrow Q_w$ to generate constraints in the language X. We can view $\Gamma \vdash e : \tau \rightsquigarrow Q_w$ as an algorithm whose inputs are Γ and e and whose outputs are τ , the type of the expression e and Q_w , the “wanted” constraint. By “wanted”, we mean that the constraint Q_w must be satisfiable for e to have type τ . The constraint Q_w is then run through a constraint solver, in an attempt to reduce Q_w to the empty constraint ϵ via simplifications and substitutions.

Constraints can also be “given” constraints. These constraints arise from user type annotations. For example, if the user has declared $\text{foo} :: \text{Coercible } a \ b \Rightarrow [a] \rightarrow [b]$, then the definition of `foo` will be type-checked under an assumption that $a \sim_{\mathbb{R}} b$. This constraint will be considered a given.

5.2.1 The solver pipeline

The solver maintains a work list of simple constraints (that is, constraints without conjunctions), with given constraints prioritised over wanted ones. It proceeds by popping the first constraint off the work list (this constraint becomes the *work item*) and then processing it through the following pipeline:

1. Types in the work item are *flattened*, whereby a type τ , possibly with type family applications, is converted into a type ζ devoid of type family applications. Such types are easier to work with in subsequent steps. Flattening τ (essentially) creates a new type variable a for every type family application $F(\bar{\sigma})$ in τ , replacing the $F(\bar{\sigma})$ with a , and then adding the $F(\bar{\sigma}) \sim_{\mathbb{N}} a$ constraint to the work list. The details are, of course, more involved; see Vytiniotis *et al.* (2011).

⁸ Specifically, we omit implication constraints, touchable variables, and the flattening substitution.

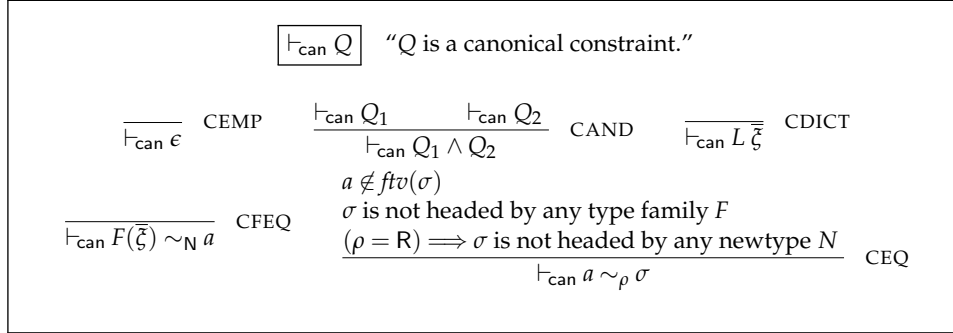


Figure 10. Canonical constraints, $\vdash_{\text{can}} Q$

2. The work item is then *canonicalised* (details are below), which reduces it to one of several simple forms. See Figure 10.
3. The work item then undergoes binary interactions with *inert* constraints, where the inert constraints are those that have already gone through this pipeline. For example, if a given inert constraint is $a \sim_N \text{Int}$, then a work item of $\text{Ord } a$ would be rewritten to Ord Int .
4. Lastly, the work item interacts with top-level axioms. This step includes type family reduction and class instance lookup.

While processing a work item, it is possible that we learn something new about a type variable, say, that a type variable b is now equal to Bool . When this happens, any inert constraint mentioning b is *kicked out* of the inert set and re-added to the work list. This step is necessary because the new knowledge about b may allow new interactions to occur.

5.2.2 Canonicalisation

The component of the solver that concerns us most is the *canonicalisation* algorithm. We write one step of this algorithm as $\text{canon}[\ell](Q_1) = Q_2$ if it succeeds, or as $\text{canon}[\ell](Q_1) = \perp$ if it fails because the constraint is unsatisfiable. The parameter ℓ is a constraint *flavour*, which can be either given (“g”) or wanted (“w”).

Canonicalisation runs the *canon* algorithm until a fixpoint is reached; the result may or may not be canonical (according to the $\vdash_{\text{can}} Q$ judgement in Figure 10). A constraint without a canonical form is not an error – perhaps later, the constraint solver will learn more and will be able to make more progress.

There are three basic canonical forms: a class constraint where all arguments are type-function-free ($\bar{\zeta}$ is a metavariable for type-function-free types), an equality between a type family application and a type variable, and an equality between a type variable and a type that is not a type family application. The empty constraint and the conjunction of canonical constraints is also canonical.

The canonicalisation algorithm must be sound with respect to constraint entailment.

Define $\text{canon}[\ell](Q_1) = Q_2 \perp$ by top-to-bottom pattern matching as follows:		
REFL	$\text{canon}[\ell](\tau \sim_R \tau)$	$= \epsilon$
NEWL	$\text{canon}[\ell](N \bar{\tau} \sim_R \sigma_2)$	$= \sigma_1[\bar{\tau}/\bar{a}] \sim_R \sigma_2$
	newtype $N \bar{a} \rightsquigarrow^* \sigma_1$, where σ_1 is not headed by a newtype	
NEWR	$\text{canon}[\ell](\sigma_1 \sim_R N \bar{\tau})$	$= \sigma_1 \sim_R \sigma_2[\bar{\tau}/\bar{a}]$
	newtype $N \bar{a} \rightsquigarrow^* \sigma_2$, where σ_2 is not headed by a newtype	
DECOMPN	$\text{canon}[\ell](H \bar{\tau} \sim_N H \bar{\sigma})$	$= \bar{\tau} \sim_N \bar{\sigma}$
DECOMP R	$\text{canon}[\ell](H \bar{\tau} \sim_R H \bar{\sigma})$	$= \bar{\tau} \sim_\rho \bar{\sigma}$
	$(H \text{ is not a newtype}) \wedge (\bar{\rho} \text{ is a prefix of } \text{roles}(H))$	
DECOMP FN	$\text{canon}[\ell](H_1 \bar{\tau} \sim_N H_2 \bar{\sigma})$	$= \perp$
DECOMP FR	$\text{canon}[\ell](H_1 \bar{\tau} \sim_R H_2 \bar{\sigma})$	$= \perp$
	$(H_1 \text{ is not a newtype}) \wedge (H_2 \text{ is not a newtype})$	
DECOMP NEW	$\text{canon}[w](N \bar{\tau} \sim_R N \bar{\sigma})$	$= \bar{\tau} \sim_\rho \bar{\sigma}$
	$(\text{no givens might match}) \wedge (\bar{\rho} \text{ is a prefix of } \text{roles}(N))$	
DECOMP D	$\text{canon}[\ell](H_1 \bar{\tau} \sim_R H_2 \bar{\sigma})$	$= H_1 \bar{\tau} \sim_R H_2 \bar{\sigma}$
APPN	$\text{canon}[\ell](\tau_1 \sigma_1 \sim_N \tau_2 \sigma_2)$	$= \tau_1 \sim_N \tau_2 \wedge \sigma_1 \sim_N \sigma_2$
APPR	$\text{canon}[\ell](\tau_1 \sigma_1 \sim_R \tau_2 \sigma_2)$	$= \tau_1 \sigma_1 \sim_R \tau_2 \sigma_2$
TVREFL	$\text{canon}[\ell](a \sim_\rho a)$	$= \epsilon$
OCCURL	$\text{canon}[\ell](a \sim_N \sigma)$	$= \perp$
	$a \in \text{ftv}(\sigma)$	
OCCUR R	$\text{canon}[\ell](\sigma \sim_N a)$	$= \perp$
	$a \in \text{ftv}(\sigma)$	
TVL	$\text{canon}[\ell](a \sim_\rho \sigma)$	$= a \sim_\rho \sigma$
TVR	$\text{canon}[\ell](\sigma \sim_\rho a)$	$= a \sim_\rho \sigma$
FAIL	$\text{canon}[\ell](\tau_1 \sim_\rho \tau_2)$	$= \perp$

Figure 11. The canonicalisation algorithm for equality constraints. These rules are explained in Section 5.3.

Property 1 (Canonicalisation soundness).

1. When $\text{canon}[g](Q_1) = Q_2$, it must be that $Q_1 \Vdash Q_2$. That is, we can derive evidence for Q_2 given evidence for Q_1 .
2. When $\text{canon}[w](Q_1) = Q_2$, it must be that $Q_2 \Vdash Q_1$. That is, by producing evidence for Q_2 we will be able to produce evidence for Q_1 .

5.3 Canonicalising equality constraints

Adding representational equalities to X requires changing the canonicalisation algorithm of $\text{OUTSIDEIN}(X)$. Indeed, with the exception of a straightforward new binary interaction (Section 5.4), this is the *only* change necessary for the constraint solver.

Figure 11 presents the portion of the algorithm that works on equality constraints of the form $\tau_1 \sim_\rho \tau_2$. A result of \perp (pronounced “failure”) means that a definite type error can be reported. For example, attempting to canonicalise $\text{Int} \sim_R \text{Bool}$ yields failure.

We describe the rules from the figure in order from top to bottom, except for REFL which we defer to Section 5.3.4.

5.3.1 Unwrapping newtypes

Rules NEWL and NEWR unwrap newtypes. They work only at the outermost level, unwrapping Age but not Maybe Age. Although not expressed in Figure 11, unwrapping newtypes happens only when the newtype’s constructor is in scope (see Section 3.1).

Both rules also unwrap eagerly, continuing to unwrap outermost newtypes until a type without an outermost newtype is found. If no such type is found (because a newtype is recursive without an intervening non-newtype), then rule NEWL (or NEWR) does not apply. This is the meaning of the \rightsquigarrow^* widget in the rules.

In the case of a recursive newtype, though, this unwrapping can diverge. For example, consider this unlikely construction:

newtype FixEither $x = \text{MkFE } (\text{Either } x \text{ (FixEither } x))$

The role of FixEither’s parameter will be inferred to be representational. Now suppose we are trying to canonicalise the wanted constraint

$$[w] \text{FixEither Age} \sim_R \text{FixEither Int}$$

Assuming the MkFE is in scope, unwrapping yields:

$$[w] \text{Either Age (FixEither Age)} \sim_R \text{Either Int (FixEither Int)}$$

Now decomposition yields $\text{Age} \sim_R \text{Int}$ (which is easily solved), and the original constraint $\text{FixEither Age} \sim_R \text{FixEither Int}$, so we are back where we started, having made no progress. In the actual implementation, a reduction counter (incremented every time a constraint is simplified and with an arbitrary, user-controllable limit) notices the loop and reports an error.

5.3.2 Decomposition of applied type constants

The DECOMP rules implement decomposition of an applied type constant. Decomposition for nominal equality is straightforward: if the two constants match, decompose (DECOMP_N); otherwise, fail (DECOMP_{FN}). For representational equality, however, the rules are more subtle. Rule DECOMP_R fires only for non-newtypes. It is easiest to understand this restriction by considering the “given” case separately from the “wanted” case.

We cannot decompose given newtype representational constraints. If the constraint in question is a given, then it would be unsound to decompose. Note rule CO_{NTH} from Figure 5, which forbids the type constant involved from being a newtype. See Section 4.2.7 for the details. Creating evidence for decomposing a given constraint of the form $H \bar{\tau} \sim_R H \bar{\sigma}$ requires using **nth**, and so we are stuck.

Wanted newtype representational constraints are tricky. We must decompose wanted newtype representational constraints, even when unwrapping does not apply. For example it happens that the abstract type `IO t` is implemented by a newtype. Haskell programmers certainly want to coerce between `(IO Int)` and `(IO Age)`. However, since `IO` is abstract, its data constructor is not visible to clients, and hence we cannot use newtype unwrapping (Section 5.3.1); so the only way forward is to decompose.

However we must tread carefully, because in certain situations it is just possible for decomposition to make a provable goal un-provable, which would compromise the completeness of type inference. Here is a contrived scenario illustrating the problem:

```
newtype ConstBool a = Mk Bool
type role ConstBool representational
```

Suppose the constructor `Mk` is not in scope. Now, consider the following constraints, where Greek letters denote unification variables:

$$\begin{array}{ll} [g] & \text{ConstBool } a \sim_R \text{ConstBool } b \quad (1) \\ [w] & \text{ConstBool } \alpha \sim_R \text{ConstBool } b \quad (2) \\ [w] & \alpha \sim_R a \quad (3) \end{array}$$

The wanted goal is certainly provable from the givens; just use (3) to substitute for α in (2), and then (2) is equal to (1). However, suppose the constraint solver happens to process (2) before (3). Because `Mk` is not in scope, `ConstBool` cannot be unwrapped. So we apply decomposition, yielding the un-satisfiable wanted goal $\alpha \sim_R b$, and hence (wrongly) report an error. This kind of incompleteness is particularly confusing to the programmer, because the goal we are trying to prove is practically equal to one of the givens.

To avoid this confusion, we do one extra check before decomposing a wanted newtype representational equality, to make sure that no givens could possibly influence the wanted constraint. This check is done by trying to unify all givens with the constraint; if any given indeed unifies, then we do not decompose. This is the informally-stated “no givens might match” side condition on the `DECOMPNEW` rule. To formalise the side condition we would need to pass to *canon* the set of (canonicalised) givens, which would clutter up Figure 11. Happily, there is no difficulty in the implementation, and the check turns out not to be as expensive as it might seem, because there are rarely many givens in practice.

There is another awkward consequence of decomposing wanted representational newtypes. Consider the `FixEither` example given in Section 5.3.1. As we saw there, canonicalising will loop if the data constructor `MkFE` is in scope. But if it is not, we will decompose to $\text{Age} \sim_R \text{Int}$, which is easily soluble. This is a situation where importing the `MkFE` constructor makes a typeable program become ill-typed, rather unfortunately. However, this seems the best we can do.

Failure and stuck cases for decomposing representational equalities. In rule `DECOMPFR`, we fail (reporting an error) when canonicalising a representational equality between two different type constants, neither of which is a newtype.

On the other hand, to account for the `ConstBool` example discussed above, `DECOMPFD` returns unchanged any remaining representational equality constraint between two applied type constants. If no earlier rule has fired, then we don't know enough about these types either to canonicalise fully or to be sure the program has a type error. These constraints will be examined again by the solver after it has learned more from other constraints.

5.3.3 Decomposition of applied type variables

Rule `APPN` decomposes an equality between two type applications, where the head of the type in the “function” position is just a type variable. (Note that the head of a nested type application must be either a type constant or a type variable; anything else would be ill-kinded.) This rule works only over nominal equality, as decomposing a representational equality of this form – say, $a\tau \sim_R b\sigma$ – is unsound, for two reasons:

- We do not know the roles on a and b . Accordingly, should we reduce to $\tau \sim_R \sigma$ or $\tau \sim_N \sigma$? It is impossible to know, especially considering that we might learn, later on during solving, the concrete value for a or b .
- Perhaps more problematic, type variables may stand in for newtypes. If we learn, say, that $a \sim_N N$ for some newtype N , then it is possible that τ and σ are unrelated, as $N\tau \sim_R b\sigma$ might be solved via unwrapping N .

Decomposing a representational equality among such type applications is not possible, but neither is this an error. We thus simply fail to canonicalise such constraints, as shown in `APPR`, which returns the same constraint it is given.

5.3.4 The reflexivity check

Rule `REFL` checks for reflexivity, succeeding with an empty constraint if the equality is reflexive. This check is needed only for representational equality constraints, as it is redundant with later checks for nominal equality: any nominal equality constraint is decomposed into its atoms, which are then checked for reflexivity. For representational equality, however, this is not the case, both because of the possibility of recursive newtypes and of impossible-to-decompose type applications.

Here are examples of these cases. Suppose we have X :

newtype $X = \text{Mk}X (\text{Int} \rightarrow X)$

Further, suppose we have these (unrelated) constraints:

$$\begin{array}{l} [w] \quad X \sim_R X \\ [w] \quad f a \sim_R f a \end{array}$$

Without the reflexivity check, canonicalising the first constraint would loop, in exactly the same way as the `FixEither` example of Section 5.3.1. Canonicalising the second constraint would simply be stuck without the reflexivity check, hitting rule `APPR` and making no progress. Programmers find it particularly frustrating if a compiler says that it is unable to prove that two syntactically-identical types are coercible, e.g. `Coercible X X!`

Note that rule `REFL` is tried first, before unwrapping newtypes, otherwise the `X` example above would loop through `NEWL/NEWR`.

5.3.5 Dealing with type variables

Rule `TVREFL` dispatches the case where we compare a type variable with itself, at either role. The *canon* algorithm then does an “occurs check” (rules `OCCURL` and `OCCURR`). The occurs check is made for nominal equality only, because an occurs-check failure for representational equality is not necessarily an error. Suppose we have $a \sim_R ba$, but then we later learn that $b \sim_N \text{Id}$, where **newtype** `Id a = Id a`. The $a \sim_R ba$ equality now becomes easily solvable.

Because of the possibility of occurs-check failures, rules `TVL` and `TVR` do not necessarily produce canonical constraints over representational equalities. Canonical type variable equality constraints must pass the occurs check, even for representational equality constraints, because they are used for substitutions. Our $a \sim_R ba$ constraint then remains non-canonical, but not otherwise harmful.

As detailed in Section 5.4, we use canonical type variable equalities as a substitution in other constraints. However, because representational equalities that fail the occurs check are not canonical, these equalities cannot be used. This can cause yet another source of incompleteness in our algorithm. Consider the following scenario:

$$\begin{array}{l} [g] \quad a \sim_R b \ a \\ [w] \quad a \text{ Int} \sim_R b \ a \text{ Int} \end{array}$$

The first, given equality cannot be canonicalised, but nor is it an outright error. When we try to solve the second, wanted equality constraint, we fail, unable to use the first for substitution.

5.3.6 Correctness of canon

Theorem (Soundness of *canon*). *The canon algorithm as presented in Figure 11 is sound, according to Property 1.*

Proof. For each rule in *canon*, it is possible to create a coercion witnessing the result from the input, and it is possible to create a coercion witnessing the input from the result, all using the coercion formation rules of Figure 5. These coercions are all straightforward to build. As the entailment relation $Q \Vdash Q$ derives from the coercion formation rules, these coercions witness the entailments we desire. \square

Note that we do not attempt to prove the algorithm complete – indeed, we know that with its treatment of recursive newtypes, type applications, and occurs-check failures, the algorithm is incomplete.

5.4 Substitution with representational equalities

Using the canonicalisation algorithm just described is nearly enough to have the `OUTSIDEIN(X)` solver work with representational equalities. The one remaining piece is to implement transitivity in the presence of assumptions. For example, we would like to be able to deduce $a_1 \sim_R a_2 \wedge a_2 \sim_R a_3 \Vdash a_1 \sim_R a_3$. This is accomplished by allowing substitution by representational equalities in representational equality constraints. (Previously, only nominal equality constraints were used for substitution.) In this case, $a_1 \sim_R a_2$ and $a_2 \sim_R a_3$ are givens. These are already in canonical form. When solving the wanted $a_1 \sim_R a_3$, we can use canonical type variable representational equality constraints to rewrite other representational equality constraints. We thus rewrite a_1 to a_2 and then a_2 to a_3 in $a_1 \sim_R a_3$. We then get $a_3 \sim_R a_3$ and are done.

This use of rewriting only works with *canonical* constraints. Transitivity is thus somewhat limited. For example, the following fails to type-check:

```
incomplete :: (Coercible (a b) (c d), Coercible (c d) (e f))
  => c d -> a b -> e f
incomplete _ = coerce
```

(The first argument is just to make c and d unambiguous.) This definition *should* be accepted, but it is not, as *canon* cannot canonicalise the givens and then discover the transitivity. We conjecture that this source of incompleteness could be overcome with more engineering, but there seems to be little incentive to add the extra complexity to the solver in this obscure case.

5.5 Properties of type inference

We have detailed our update to the canonicalisation algorithm and have given a sketch of the overall type inference algorithm. Here we informally discuss some of the attributes of type inference in Haskell augmented with `Coercible`.

Incompleteness Our type inference algorithm remains incomplete, following on from the known incompleteness of `OUTSIDEIN(X)` (e.g., see Vytiniotis *et al.* (2011), Section 6.5). However, our approach specifically toward representational equality adds new forms of incompleteness. We saw examples in Sections 5.3.1, 5.3.2, and 5.3.5. We have considered ways to improve the algorithm to handle these cases and believe there is opportunity for such improvement. However, we have been unable to find a principled approach that would be a clear improvement over the algorithm we present here, in terms of realistic programs that are newly accepted and in terms of engineering effort and complexity. We expect that this algorithm will evolve as more users make practical use of `Coercible` and identify places where the algorithm’s incompleteness is a concrete barrier to progress.

Decidability Though we have not formally proved it, we conjecture that solving for representational equality is undecidable. Determining whether or not two types are representationally equal is essentially an equality check on equirecursive higher-kinded types, which would appear to subsume equivalence of terms in a λ -calculus with a fix operator. This conjectured lack of decidability, if indeed true, prevents us from ever writing a complete algorithm.

One still might ask whether the incomplete algorithm we have written is guaranteed to terminate. As discussed in Section 5.3.1, our implementation must use a counter to prevent divergence in the presence of recursive newtypes. Accordingly our algorithm, as stated, is not guaranteed to terminate. However, we find that this is not problematic in practice, both because of the presence of a counter to detect runaway recursion and the fact that non-termination is believed to happen only when the user requires us to solve an undecidable problem: representational equality among recursive newtypes. When a user asks for such a thing, we are not terribly ashamed when we take forever in delivering it.

6 Reflection and discussion

This section discusses some opportunities and choices that arose in the course of our work on safe coercions.

6.1 Generalized Newtype Deriving done right

As mentioned before, **newtype** is a great tool to make programs more likely to be correct, by having the type checker enforce certain invariants or abstractions. But newtypes can also lead to tedious boilerplate. Assume the programmer needs an instance of the type class `Monoid` for her type `HTML`. The underlying type `String` already comes with a suitable instance for `Monoid`. Nevertheless, she has to write quite a bit of code to convert that instance into one for `HTML`:

```
instance Monoid HTML where
  mempty = MkHTML mempty
  mappend (MkHTML a) (MkHTML b) = MkHTML (mappend a b)
  mconcat xs = MkHTML (mconcat (map unHTML xs))
```

Note that this definition is not only verbose, but also non-trivial, as invocations of `MkHTML` and `unHTML` have to be put in the right places, possibly via some higher order functions like `map` – all just to say “just use the underlying instance”!

This task is greatly simplified with `Coercible`: Instead of wrapping and unwrapping arguments and results, she can directly coerce the method of the base type’s instance itself:

```
instance Monoid HTML where
  mempty = coerce (mempty :: String)
  mappend = coerce (mappend :: String → String → String)
  mconcat = coerce (mconcat :: [String] → String)
```

```

newtype Id1 a = MkId1 a
newtype Id2 a = MkId2 (Id1 a)
  deriving (UnsafeCast b)

type family Discern a b
type instance Discern (Id1 a) b = a
type instance Discern (Id2 a) b = b

class UnsafeCast to from where
  unsafe :: from → Discern from to

instance UnsafeCast b (Id1 a) where
  unsafe (MkId1 x) = x

unsafeCoerce :: a → b
unsafeCoerce x = unsafe (MkId2 (MkId1 x))

```

Figure 12. The above implementation of `unsafeCoerce` compiles (with appropriate flags) in GHC 7.6.3 but does not in GHC 7.8.1.

The code is pure boilerplate: apply `coerce` to the method, instantiated at the base type by a type signature. And because it is boilerplate, the compiler can do it for her; all she has to do is to declare which instances of the base type should be lifted to the new type by listing them in the **deriving** clause:

```

newtype HTML = MkHTML String
  deriving Monoid

```

This is not a new feature: GHC has provided this *Generalized Newtype Deriving* (GND) for many years. But, the implementation was “magic” – GND would produce code that a user could not write herself. Now, the feature can be explained easily and fully via `coerce`.

Furthermore, GND was previously unsound (Weirich *et al.*, 2011). When combined with other extensions of GHC, such as type families (Chakravarty *et al.*, 2005a; Chakravarty *et al.*, 2005b) or GADTs (Cheney & Hinze, 2003), GND could be exploited to completely break the type system: Figure 12 shows how this notorious bug can allow any type to be coerced to any other. The clause “**deriving** (UnsafeCast b)” is the bogus use of GND, and now will generate the instance

```

instance UnsafeCast b (Id2 a) where
  unsafe = coerce (unsafe :: Id1 a → Discern (Id1 a) b)

```

which will rightly be rejected because `Discern`’s first parameter has a nominal role. Indeed, preventing abuse of GND was the entire subject of (Weirich *et al.*, 2011).

Similarly, it was possible to use GND to break invariants of abstract data types. The addition of `coerce` makes it yet easier to break such abstractions. As discussed in Section 3.1, these abuses can now be prevented via role annotations.

6.2 Coercible and rewrite rules

What if a client of module `Html` writes this?

```
....(map unHTML hs)...
```

She cannot use `coerce` because `HTML` is an abstract type, so the type system would (rightly) reject an attempt to use `coerce` (Section 3.1). However, since `HTML` is a newtype, one might hope that GHC's optimiser would transform `(map unHTML)` to `coerce`. The optimiser must respect type soundness, but (by design) it does not respect abstraction boundaries: dissolving abstractions is one key to high performance.

The correctness of transforming `(map unHTML)` to `coerce` depends on a theorem about `map`, which a compiler can hardly be expected to identify and prove all by itself. Fortunately GHC already comes with a mechanism that allows a library author to specify *rewrite rules* for their code (Peyton Jones *et al.*, 2001). The author takes the proof obligation that the rewrite is semantics-preserving, while GHC simply applies the rewrite whenever possible. In this case the programmer could write

```
{-# RULES "map/co" map coerce = coerce #-}
```

In our example, the programmer wrote `(map unHTML)`. The definition `unHTML` in module `Html` does not mention `coerce`, but both produce the same System FC code (a cast). So via cross-module inlining (more dissolution of abstraction boundaries) `unHTML` will be inlined, transforming the call to the equivalent of `(map coerce)`, and that in turn fires the rewrite rule. Indeed even a nested call like `map (map unHTML)` will also be turned into a single call of `coerce` by this same process applied twice.

The bottom line is this: the author of a map-like function `someMap` can accompany `someMap` with a `RULE`, and thereby optimise calls of `someMap` that do nothing into a simple call to `coerce`.

Would it be sufficient to expose only this mechanism on the Haskell source level to achieve our goals, without introducing a user-visible `coerce` function and `Coercible` constraint? No, for a number of reasons:

- Rewrite rules are but an optimisation, e.g. they are only applied when the compiler is run with the `-O` command line option. The `coerce` function is always a zero-cost function.
- Rewrite rules are rather difficult to master. It takes practice to predict when they will fire, and it is tricky to investigate the reasons in case they do not fire. The programmer would hence have a hard time to guarantee that a certain piece of code is indeed compiled down to a zero-cost conversion.
- When converting between two complex types, the programmer would have to puzzle together a possible large number of map-like functions, **newtype** constructors and destructors to build a rather complicate chunk of code, just to let the compiler optimise it away. With `coerce`, he writes just that and the compiler does the rest of the work.

- Moreover, some of these map-like functions might not be exported, might not have rules attached to them or might not even exist in the first place, barring the user from implementing the desired conversion by foot.

6.3 Syntax for role annotations

Recall the Map example from Section 3.1, and its role annotation:

```
data Map k v = Leaf | Node k v (Map k v) (Map k v)
type role Map nominal representational
```

This is only one possible concrete syntax for role annotations, and we explored a number of others. In doing so we identified the following design criteria:

1. Role annotations must be optional. Otherwise, all existing code would be broken.
2. Role annotations should be succinct.
3. Role annotations will be a relatively obscure feature, and therefore should be searchable should a user come across one.
4. Code with role annotations should compile with older versions of GHC, easing migration to the first version of GHC supporting roles (GHC 7.8).
5. Role annotations should not be specified in a pragma; pragmas are meant to be reserved for implementation details (e.g., optimising), and roles are a type system feature.
6. Role annotations should be easy to refactor as a data type evolves.
7. Code is read much more often than it is written; favour readability over concision.

Our chosen syntax, with **type role ...**, satisfies criteria (1), (3), (5), and (7), at the cost of some others. In particular, this choice is not backward compatible. A role annotation fails to parse in earlier versions of GHC. However, GHC supports C-style preprocessor directives, so library authors can selectively include role annotations using preprocessor directives. The fact that the annotations are standalone means they can be grouped under one set of directives instead of sprinkled throughout the source file. Note that this syntax is very easy to search for and the written-out nature of the roles makes them readable, if not so concise. Breitner *et al.* (2014b) discusses alternatives to this syntax in Appendix B.1.

6.4 The role of role inference

Why did we add role inference to GHC, assigning the most permissive role to type constructors by default?

We did not have to design the language in this way. We could have required programmers to annotate the roles of every type, which GHC would check for consistency. However, in this case the burden on programmers seems drastic and migration to this system overwhelming, requiring all existing data type declarations to be annotated with roles.

Alternatively, we could specify that all unannotated roles default to nominal (thus removing the need for role inference). According to the specification of Figure 7, it is always sound to assign the nominal role to all parameters of a type constructor H . This choice would lead to greater abstraction safety by default. For example, the implementor of `Map` would not need to add a role annotation to guarantee abstraction.

However, we choose to use the most permissive roles by default for several reasons. First, for convenience: this choice increases the availability of `coerce` (as only those types with annotations would be `Coercible` otherwise), and it supports backward compatibility with the Generalized Newtype Deriving (GND) feature (see Section 6.1).

Furthermore, role inference also means that the majority of programmers do not need to learn about roles nor need to add role annotations. Users of `coerce` will need to consider roles, as will library implementors who use class-based invariants (see Section 3.1). Other users are unaffected by roles and will not be burdened by them.

Our choices in the design of the role system has generated vigorous debate.⁹ This discussion is healthy for the Haskell community. The difficulty with abstraction is not new: with GND, it has always been possible to lift coercions through data types, potentially violating their class-based invariants. The features described in this paper make this subversion both more convenient (through the use of `coerce`) and, more importantly, now preventable (through the use of role annotations).

6.5 Roles in Practice

We have described a mechanism to allow safe coercions among distinct types, and we have reimplemented GHC's previously unsafe Generalized Newtype Deriving extension in terms of these safe coercions. Naturally, this change causes some code that was previously accepted to be rejected. Given that Haskell has a large user base and a good deal of production code, how does this change affect the community?

Advance testing During the development of this feature, we tested it against several popular Haskell packages available through Hackage, an online Haskell open-source distribution site. These tests were all encouraging and did not find any instances of hard-to-repair code in the wild.

Compiling all of Hackage As of 30 September 2013, 3,234 packages on Hackage compiled with GHC 7.6.3, the last released version without roles. The development version of GHC at that time included roles. A total of only four packages

⁹ To read some of this debate, see the thread beginning with this post: <http://www.haskell.org/pipermail/libraries/2014-March/022321.html>

failed to compile directly due to GND failure.¹⁰ Of these, three of the failures were legitimate – the use of GND was indeed unsafe. For example, one case involved coercing a type variable passed into a type family; the author implicitly assumed that a newtype and its representation type were always considered equivalent with respect to the type family. Only one package – *acme-schoenfinkel* – failed to compile because of the gap in expressiveness between the roles in Weirich *et al.* (2011) and those here. No other Hackage package depends on this one, indicating it is not a key part of the Haskell open-source fabric. The example in Section 8.1.1 is along similar lines to the failure observed here.

These data were gathered almost two months after the implementation of roles was pushed into the development version of GHC, so active maintainers may have made changes to their packages before the study took place. Indeed, we are aware of a few packages that needed manual updates. In these cases, instances previously derived using GND had to be written by hand, but quite straightforwardly.

Rewrite rules Since GHC 7.10, the rewrite rule *"map/co"* (Section 6.2) has been added to the standard library, and indeed, it does fire: We analysed 1,077 packages¹¹. In 64 of these packages, the *"map/co"* rule fired and eliminated a total of 272 calls to *map* (out of 13,991 calls that were not already dissolved by list fusion).

7 Related work

Prior work discusses the relationship between roles in FC and languages with generativity and abstraction, type-indexed constructs, and universes in dependent type theory. We do not repeat that discussion here. Instead we use this section to clarify the relationship between this paper and Weirich *et al.* (2011), as well as make connections to other systems.

7.1 Prior version of roles

The idea of *roles* was initially developed in Weirich *et al.* (2011) as a solution to the Generalized Newtype Deriving problem. That work introduces the equality relations \sim_R and \sim_N (called “type equality” and “code equality” resp. in Weirich *et al.* (2011)). However, the system presented in Weirich *et al.* (2011) was quite invasive: it required annotating every sub-tree of every kind with a role. Kinds in GHC are already quite complicated because of kind polymorphism, and a new form of role-annotated kinds would be more complex still.

In this paper, we present a substantially simplified version of the roles system of Weirich *et al.* (2011), requiring role information only on the parameters to data

¹⁰ These data come from Bryan O’Sullivan’s work, described here: <http://www.haskell.org/pipermail/ghc-devs/2013-September/002693.html> That posting includes 3 additional GND failures; these were due to an implementation bug, since fixed.

¹¹ Stackage nightly 2015-05-21, excluding two packages with non-Haskell dependencies that were not fulfilled on our test system

types. Our new design keeps roles and kinds modularly separate, so that roles can be handled almost entirely separately (both intellectually and in the implementation) from kinds. The key simplification is to “assume the worst” about higher-kinded parameters, by assuming that their arguments are all nominal. In exchange we give up some expressiveness; specifically, we give up the ability to abstract over type constructors with non-nominal argument roles (see Section 8.1).

Furthermore, the observation that it is sound to “assume the worst” and use parameterised types with less permissive roles opens the door to role annotations. In this work, programmers are allowed to deliberately specify less permissive roles, giving them the ability to preserve type abstractions.

Surprisingly, this flexibility means that our version of roles actually *increases* expressiveness compared to Weirich *et al.* (2011) in some places. In Weirich *et al.* (2011) a role is part of a type’s kind, so a type expecting a higher-kinded argument (such as `Monad`) would also have to specify the roles expected by its argument. Therefore if `Monad` is applicable to `Maybe`, it would not also be applicable to a type `T` whose parameter has a nominal role. In the current work, however, there is no problem because `Maybe` and `T` have the same kind.

Besides the simplification discussed above, this paper makes two other changes to the specification of roles presented in Weirich *et al.* (2011).

- The treatment of the phantom role is entirely novel; the rule `CO_PHANTOM` has no analogue in prior work.
- The coercion formation rules (Figure 5) are refactored so that the role on the coercion is an *output* of the (syntax-directed) judgement instead of an input. This is motivated by the implementation (which does not know the role at which coercions should be checked) and requires the addition of the `CO_SUB` rule.

There are, of course, other minor differences between this system and Weirich *et al.* (2011) in keeping with the evolution of System FC. The main significant change, unrelated to roles, is the re-introduction of **left** and **right** coercions; see Section 4.2.7.

One important non-difference relates to the linear-pattern requirement. Section 4.4 describes that our language is restricted to have only *linear* patterns in its type families. (GHC, on the other hand, allows non-linear patterns as well.) This restriction exists in the language in Weirich *et al.* (2011) as well. Section 4.2.2 of Weirich *et al.* (2011) defines so-called Good contexts as having certain properties. Condition 1 in this definition subtly implies that all type families have linear patterns – if a type family had a non-linear pattern, it would be impossible, in general, to establish this condition. The fact that the definition of Good implies linear patterns came as a surprise, further explored in Eisenberg *et al.* (2014). The language described in the present paper clarifies this restriction, but it is not a new restriction.

Finally, because this system has been implemented in GHC, this paper discusses more details related to compilation from source Haskell. In particular, the role inference algorithm of Section 4.6 is a new contribution of this work.

7.2 Prior version of Coercible

This paper describes Coercible as implemented in GHC 7.10, using a dedicated solver in the type checker to handle representational equality constraints (Coercible) as well as nominal equality constraints (\sim). This approach differs from the initial design that was shipped with GHC 7.8 and discussed in an earlier version of this work (Breitner *et al.*, 2014a), where Coercible was presented as a type class instead of a special constraint.

In particular, our prior work explains the solving of Coercible constraints in terms of type class instances. The motivation was to make it possible for the programmer to predict and understand the behaviour of the compiler without special knowledge, assuming she is aware of type classes.

Unfortunately, that approach had a few drawbacks. Although it was sold as behaving “like a normal type class”, that was never fully true, and the solver treated Coercible special in a few cases:

- It would refrain from building recursive evidence. Recursive evidence is common and useful with type classes, but for Coercible it would simply cause the program to loop when executed, so we gave a compile time error instead.
- It allowed constraints of the form `Coercible (forall a. s) (forall a. t)` which are forbidden for type classes, but required here to deal with newtypes such as `newtype Sel = MkSel (forall a. [a] → a)`.
- While type class instances are always exported and unconditionally visible, the visibility of the newtype unwrapping instance depends on whether the constructor is in scope.

In the end we found it clearer to stop pretending Coercible is a type class and honestly call it a constraint of its own right, with its own rules and its own solver. This also made the feature more powerful, as the instance-based approach is not able to decompose given Coercible constraints (Section 2.7).

7.3 OCaml and variance annotations

The interactions between sub-typing, type abstraction, and various type system extensions such as GADTs and parameter constraints also appear in the OCaml language. In that context, *variance annotations* act like roles; they ensure that sub-type coercions between compatible types are safe. For example, the type `α list` of immutable lists is covariant in the parameter `α` : if `$\sigma \leq \tau$` then `σ list \leq τ list`. Variances form a lattice, with *invariant*, the most restrictive, at the bottom; *covariant* and *contravariant* incomparable; and *bivariant* at the top, allowing sub-typing in both directions. It is tempting to identify invariant with nominal and bivariant with phantom, but the exact connection is unclear. Scherer and Rémy (2013) show that GADT parameters are not always invariant.

Exploration of the interactions between type abstraction, GADTs, and other features have recently revealed a soundness issue in OCaml¹² that has been con-

¹² <http://caml.inria.fr/mantis/view.php?id=5985>

firmed to date back several years. Garrigue (2013) discusses these issues. His proposed solution is to “assume that nothing is known about abstract types when they are used in parameter constraints and GADT return types” – akin to assigning nominal roles. However, this solution is too conservative, and in practice the OCaml 4.01 compiler relies on no fewer than *six* flags to describe the variance of type parameters. However, lacking anything equivalent to Core and its tractable metatheory, the OCaml developers cannot demonstrate the soundness of their solution in the way that we have done here.

What is clear, however, is that generative type abstraction interacts in interesting and non-trivial ways with type equality and sub-typing. Roles and type-safe coercion solve an immediate practical problem in Haskell, but we believe that the ideas have broader applicability in advanced type systems.

8 Future directions

As of the date of writing (June 2015), roles seem not to have caused an undue burden to the community. The first release candidate for GHC 7.8 was released on 3 February 2014, followed by the full release on 9 April, and package authors had been updating their work to be compatible for some time. The authors of this paper are unaware of any major problems that Haskellers have had in updating existing code. However, two problems have been identified: the need for roles to work in higher-order scenarios, and the need for a better interaction between roles and Safe Haskell (Terei *et al.*, 2012). We also review some proposed expansions of the roles feature to more exotic Haskell constructs.

8.1 Roles for higher-order types

Some users wish to use roles in scenarios that are currently beyond the ability of roles to express. We focus on one such scenario, as it is representative of all examples we have seen, including the package that did not compile when testing all of Hackage (Section 6.5).

8.1.1 Adding *join* to *Monad*

Imagine adding the *join* method to the *Monad* class, as follows:

```
class Monad m where
  ...
  join :: forall a. m (m a) → m a
```

With this definition, GND would still work in many cases. For example, if we define

```
newtype M a = MkM (Maybe a)
deriving Monad
```

GND will work without a problem. We would need to show `Coercible (Maybe (Maybe a) → Maybe a) (M (M a) → M a)`, which is straightforward.

More complicated constructions run into trouble, though. Take this definition, written to restrict a monad’s interface:

```
newtype Restr m a = MkRestr (m a)
deriving Monad
```

To perform GND in this scenario, we must prove `Coercible (m (m a) → m a) (Restr m (Restr m a) → Restr m a)`. In solving for this constraint, we eventually simplify to `Coercible (m (m a)) (m (Restr m a))`. At this point, we are stuck, because we do not have any information about the role of `m`’s parameter, so we must assume it is nominal. The GND feature is thus not available here. Similar problems arise when trying to use GND on monad transformers, a relatively common idiom.

How would this scenario play out under the system proposed in Weirich *et al.* (2011)? This particular problem wouldn’t exist – `m`’s kind could have the right roles – but a different problem would. A type’s kind also stores its roles in Weirich *et al.* (2011). This means that `Monad` instances could be defined only for types that expect a representational parameter. Yet, it is sometimes convenient to define a `Monad` instance for a data type whose parameter is properly assigned a nominal role. The fact that the system described in this paper can accept `Monad` instances both for types with representational parameters and nominal parameters is a direct consequence of the *Role assignment narrowing* theorem (Section 4.3), which does not hold of the system in Weirich *et al.* (2011).

8.1.2 Implication constraints

Looking forward, there is a proposal to indeed add `join` to `Monad`, and so we want to be able to allow the use of GND on this enhanced `Monad` class. One promising approach to this problem is to allow *user-specified implication constraints*.

Continuing the example from above, imagine we could write the following:

```
deriving instance (Monad m, forall a b. Coercible a b ⇒ Coercible (m a) (m b))
⇒ Monad (Restr m)
```

When we are trying to simplify `Coercible (m (m a)) (m (Restr m a))`, we see that this constraint can be solved if `Coercible (m a) (Restr m a)`, and so we simplify. This last constraint is easy to solve via the definition of `Restr`, and so we succeed.

The constraint `forall a b. Coercible a b ⇒ Coercible (m a) (m b)` is an *implication constraint* (Hinze & Peyton Jones, 2000), saying that `Coercible (m a) (m b)` holds whenever `Coercible a b` holds, for universally-quantified type variables `a` and `b`. These constraints do not currently exist in Haskell, but users have wanted them for some time.¹³ With such constraints, it would seem that we can effectively assign roles to parameters of type variables, much like we already assign roles to parameters of type constants. For example, the implication constraint above gives

¹³ See <https://ghc.haskell.org/trac/ghc/ticket/2256>, which was created in 2008.

the parameter to `m` a representational role. This role assignment is precisely what is needed to use GND with `Monad` and `Restr`.

The details of this have yet to be fully worked out, but we believe that the implementation could be straightforward, given that GHC already deals with internal implication constraints, derived from type-checking GADT pattern-matches.

8.2 Roles and Safe Haskell

Safe Haskell (Terei *et al.*, 2012) is a subset of Haskell known to have additional safety properties. Safe Haskell excludes constructs such as the infamous functions `unsafeCoerce` and `unsafePerformIO`, as these can be used to subvert the type system. It also excludes Template Haskell (Sheard & Peyton Jones, 2002), as that feature can look up type definitions and thus break abstraction. See the original paper for the details.

One of the consequences of the unsoundness of earlier versions of GND is that the feature was (quite rightly) excluded from the Safe Haskell subset. However, even with roles and GND written in terms of `coerce`, the feature *still* does not meet the Safe Haskell criteria. At issue is preserving datatype abstraction.

We describe in Section 3.1 that we allow coercions to happen even on data types for which the constructors are not available, such as `Map`. However, this violates Safe Haskell’s promise that no abstraction barrier is broken through. To rectify this problem, GHC could use a more stringent check when satisfying a `Coercible` constraint when compiling in Safe mode, requiring all constructors of all data types to be coerced under to be visible. This means, essentially, traversing the entire tree of data type definitions, making sure all constructors of all data types, recursively, are available. We did not go this path not only because of the performance penalty of potentially having to load further interface files, but also as it would require users to import many constructors that remain unmentioned in their code, just to satisfy this requirement. We continue to look for a better solution to this problem; for some ideas, the reader is encouraged to consult <https://ghc.haskell.org/trac/ghc/wiki/SafeRoles>.

8.3 Conservativity of roles

8.3.1 Roles are coarse-grained

The system we describe has exactly three roles. However, by having only three roles, we have created a rather coarse-grained classification system, and a more fine-grained system is imaginable.

For example, consider the following definitions:

```
type family F a
type instance F Int = Char
type instance F Bool = Char
type instance F [a] = ()
```

— `F`’s parameter has a nominal role, as do all type family parameters

```
data Bar a = MkBar (F a)
type role Bar nominal
```

Is it safe to coerce a `Bar Int` to a `Bar Bool`? Unravelling definitions, we see that this is so. Yet, coercing `Bar Int` to `Bar [Double]` is clearly not safe. GHC assigns a nominal role to the parameter of `Bar`, but this choice of role eliminates the possibility of the `Bar Int` to `Bar Bool` coercion.

In order to express this, we would need to assign `Bar`'s parameter `a` a role that corresponds to the equivalence relation generated by the type family `F`, i.e. the relation between types that are mapped to representationally equal types by `F`.

To expand this example, consider `T`, which refers back to the `F` above:

```
data T a = MkT (F a)
```

Values of type `T a` share a representation with those of type `T b` precisely when `F a` is coercible to `F b`. With an expanded language for roles, we can imagine setting `a`'s role to say that any changes to `a` must respect the definition of `F` in order for `T a` to be coercible to `T b`.

Going down this route would turn our current three-role system into one with a very rich structure of equivalence relations, indexed by which type family (or even type families) are to be respected.

We could similarly imagine expressing the relation that certain type class instances are to be respected; this could allow the coercion of a `Map Int v` to a `Map Age v` precisely when `Int`'s and `Age`'s `Ord` instances correspond.

8.3.2 Equality does not propagate roles

What role should be assigned to a parameter with an equality constraint involving a phantom? According to the rules in our formalism, such a parameter would get a nominal role. Consider the following type:

```
data T a b where
  MkT :: (a ~ b) => a -> T a b
```

Role inference assigns both parameters to have nominal roles.

But this is stricter than necessary, as it disallows certain coercions. Inspection of the type definition shows us that the second parameter, `b`, is used only in the equality constraint with `a`. Additionally, ignoring the equality constraint for a moment, `a` is used only representationally. So we can conclude that `T τ τ` has the same run-time representation as `T σ σ` whenever `τ` has the same run-time representation as `σ`. Yet, the role mechanism is not expressive enough to prove this.

8.4 Extending roles to families

8.4.1 Roles on type and data families

In GHC today, all type and data family parameters have nominal roles, because a type or data family can pattern-match on its parameters. For example:

```
type family TF a
type instance TF Int = Double
type instance TF Age = Char
```

Clearly, `TF Int` is not representationally equal to `TF Age`.

Yet, it would be sensible to extend the idea of roles to type and data families. A family with a non-nominal parameter would need extra checks on its instance declarations, to make sure that they are compatible with the choice of roles. For example:

```
type role If nominal representational representational
type family If (a :: Bool) b c
type instance If True b c = b
type instance If False b c = c
```

The above definition, though not accepted by our implementation, is perfectly type safe. Note that a representational parameter must not be matched on and must not be used in a nominal context on the right-hand side. The only barrier to implementing this is the extra complexity for the GHC maintainers and the extra complexity in the language. If a compelling use case for this comes up, we will likely add the feature.

8.4.2 Roles on data family instances

Roles on data families follow the same arguments as above. However, we can identify a separate issue involving roles on data family instances, which are, of course, data types. For example:

```
data family DF a
data instance DF (b, Int) = MkDF (Maybe b)
```

Data family instances are internally desugared into something resembling a type family instance and a fresh data type declaration, somewhat like this:

```
type family DF a
type instance DF (b, Int) = DFPairIntInstance b
data DFPairIntInstance c = MkDF (Maybe c)
```

Here, it is apparent that `c` can be assigned a representational role, even while we require a nominal role for `a`.

Role inference for data family instances is not currently implemented, though it would seem to take only the will to do so.¹⁴ Instead, all type variables in a data family instance are assigned nominal roles. Why? Essentially because there is no way of writing a role annotation for data family instances. Without the ability to write role annotations, library writers would be unable to enforce abstraction on these, and so it is safer just to default these (somewhat uncommon) parameters to have nominal roles.

¹⁴ This task is tracked at <https://ghc.haskell.org/trac/ghc/ticket/8177>.

```

data Maybe a = Nothing | Just a
data Option a = None | Some a
data Few a = Zero | One a | Two a a

maybe2option :: Maybe a → Option a
maybe2option Nothing = None
maybe2option (Just x) = Some x

maybe2few :: Maybe a → Few a
maybe2few Nothing = Zero
maybe2few (Just x) = One x

```

Figure 13. Data type conversions

8.5 What else is there to coerce?

The starting point of this work was the observation that there exist expressions, such as `map MkHTML`, which change the types, but not the representation of their arguments. We built a system to express and use this in a type-safe manner.

But `Coercible` and `coerce` currently cannot be used in all such situations. Consider the data types `Maybe a` and `Option a` in Figure 13, which have – up to the names of the constructors – identical definitions. For a given compiler, it may be the case that a value `m :: Maybe a` has precisely the same representation as its counterpart `(maybe2option m) :: Option a`. If this is indeed the case, we could replace `maybe2option` with a zero-cost coercion. We expect that it would be possible to extend our system to allow for `Coercible (Maybe a) (Option a)`, in the situations where the compiler makes the two indistinguishable.

Generic programming techniques (Rodriguez *et al.*, 2008) could, if tailored around this feature, gain performance boosts if the translation between the concrete to the generic representation no longer incurs a runtime cost.

One could go even further, however. The conversion function `maybe2few` in the same Figure may also (depending on the compiler) be operationally the identity. For example, if the first constructor is tagged 1, the second is tagged 2, and so on, then `(Just x)` and `(One x)` would have the same representation. However, the situation is now asymmetrical: we may be able to convert from `Maybe a` to `Few a` for free, but the reverse is certainly not true, because the value might use the constructor `Two`.

Such uni-directional version of `Coercible` amounts to *explicit inclusive subtyping* and is more complicated than our current symmetric system: For example, the lifting rule would have to take variance into account: For a type constructor `T`, does `Coercible (T a) (T b)` require `Coercible a b`, or `Coercible b a`, or both, or neither? Furthermore, we would have to adapt our internal language, `FC`, to work with explicit subtyping proofs (Crary, 2000; Rémy & Yakobowski, 2010; Cretin & Rémy, 2012).

9 Conclusion

Our focus has been on Haskell, for the sake of concreteness, but we believe that this work is important beyond the Haskell community. Any language that offers *both* generative type abstraction *and* type-level type discrimination must deal with their interaction, and those interactions are extremely subtle. We have described one sound and tractable way to combine the two, including the source language changes, type inference, core calculus, and metatheory. In doing so we have given a concrete foundation for others to build upon.

Acknowledgements

Thanks to Antal Spector-Zabusky for contributing to this version of FC; and to Edward Kmett and Dimitrios Vytiniotis for discussion and feedback. This material is based upon work supported by the National Science Foundation under grant nos. CCF-1116620 and CCF-1319880. The first author was supported by the Deutsche Telekom Stiftung and the second author was supported by a Microsoft Research PhD Fellowship. This paper is typeset with the help of ott (Sewell *et al.*, 2010).

References

- Breitner, Joachim, Eisenberg, Richard A., Peyton Jones, Simon, & Weirich, Stephanie. (2014a). Safe zero-cost coercions for Haskell. *ICFP*. ACM.
- Breitner, Joachim, Eisenberg, Richard A., Peyton Jones, Simon, & Weirich, Stephanie. (2014b). *Safe zero-cost coercions for Haskell (extended version)*. Tech. rept. MS-CIS-14-07. University of Pennsylvania.
- Chakravarty, Manuel M. T., Keller, Gabriele, & Peyton Jones, Simon. (2005a). Associated type synonyms. *ICFP*. ACM.
- Chakravarty, Manuel M. T., Keller, Gabriele, Peyton Jones, Simon, & Marlow, Simon. (2005b). Associated types with class. *POPL*. ACM.
- Cheney, James, & Hinze, Ralf. (2003). *First-class phantom types*. Tech. rept. Cornell University.
- Crary, Karl. (2000). Typed compilation of inclusive subtyping. *ICFP*. ACM.
- Cretin, Julien, & Rémy, Didier. (2012). On the power of coercion abstraction. *POPL*. ACM.
- Eisenberg, Richard A., Vytiniotis, Dimitrios, Peyton Jones, Simon, & Weirich, Stephanie. (2014). Closed type families with overlapping equations. *POPL*. ACM.
- Garrigue, Jacques. 2013 (Sept.). *On variance, injectivity, and abstraction*. OCaml Meeting, Boston.
- Hall, Cordelia V., Hammond, Kevin, Peyton Jones, Simon L., & Wadler, Philip L. (1996). Type classes in Haskell. *Acm trans. program. lang. syst.*, **18**(2).
- Hinze, Ralf, & Peyton Jones, Simon. (2000). Derivable type classes. *Haskell Workshop*.
- Marlow, Simon (editor). (2010). *Haskell 2010 language report*.
- Milner, Robin, Tofte, Mads, Harper, Robert, & MacQueen, David. (1997). *The definition of Standard ML (revised)*.
- Peyton Jones, Simon, Tolmach, Andrew, & Hoare, Tony. (2001). Playing by the rules: rewriting as a practical optimisation technique in GHC. *Haskell Workshop*.
- Pottier, François, & Rémy, Didier. (2005). The essence of ML type inference. *Chap. 10 of: Pierce, Benjamin C. (ed), Advanced topics in types and programming languages*. MIT Press.

- Rémy, Didier, & Yakobowski, Boris. (2010). A Church-style intermediate language for MLF. *Pages 24–39 of: Blume, Matthias, Kobayashi, Naoki, & Vidal, Germán (eds), Functional and logic programming*. LNCS, vol. 6009. Springer.
- Rodriguez, Alexey, Jeuring, Johan, Jansson, Patrik, Gerdes, Alex, Kiselyov, Oleg, & Oliveira, Bruno C. d. S. (2008). Comparing libraries for generic programming in Haskell. *Haskell Symposium*. ACM.
- Scherer, Gabriel, & Rémy, Didier. (2013). GADTs meet subtyping. *ESOP*. LNCS, vol. 7792. Springer.
- Sewell, Peter, Zappa Nardelli, Francesco, Owens, Scott, Peskine, Gilles, Ridge, Thomas, Sarkar, Susmit, & Strniša, Rok. (2010). Ott: Effective tool support for the working semanticist. *Journal of functional programming*, **20**(1).
- Sheard, Tim, & Peyton Jones, Simon. (2002). Template meta-programming for Haskell. *Haskell Workshop*. ACM.
- Sulzmann, Martin, Chakravarty, Manuel M. T., Peyton Jones, Simon, & Donnelly, Kevin. (2007). System F with type equality coercions. *TLDI*. ACM.
- Terei, David, Marlow, Simon, Peyton Jones, Simon, & Mazières, David. (2012). Safe Haskell. *Haskell Symposium*. ACM.
- Vytiniotis, Dimitrios, Peyton Jones, Simon, Schrijvers, Tom, & Sulzmann, Martin. (2011). OutsideIn(X) modular type inference with local assumptions. *Journal of functional programming*, **21**(4-5).
- Weirich, Stephanie, Vytiniotis, Dimitrios, Peyton Jones, Simon, & Zdancewic, Steve. (2011). Generative type abstraction and type-level computation. *POPL*. ACM.
- Yorgey, Brent A., Weirich, Stephanie, Cretin, Julien, Peyton Jones, Simon, Vytiniotis, Dimitrios, & Magalhães, José Pedro. (2012). Giving Haskell a promotion. *TLDI*. ACM.

A System FC, in full

Throughout this entire proof of type safety, any omitted proof is by (perhaps mutual) straightforward induction on the relevant derivations.

As usual, all definitions and proofs are only up to α -equivalence. If there is a name clash, assume a variable renaming to a fresh variable.

A.1 The remainder of the grammar

Φ	::=	$[\overline{a:\kappa}].\tau \sim_\rho \sigma$	axiom types
e	::=		expressions
		x	variable
		$\lambda x:\tau.e$	value abstraction
		$\Lambda a:\kappa.e$	type abstraction
		$\lambda c:\phi.e$	coercion abstraction
		K	data constructor
		$e_1 e_2$	application
		$e \tau$	type application
		$e \gamma$	coercion application
		case $_\tau e$ of \overline{alt}	pattern match
		$e \triangleright \gamma$	cast
		contra $\gamma \tau$	absurdity
v	::=		expression values
		$\lambda x:\tau.e$	value abstraction
		$\Lambda a:\kappa.v$	type abstraction
		$\lambda c:\phi.e$	coercion abstraction
		$K \overline{\tau} \overline{\gamma} \overline{e}$	applied data constructor
alt	::=	$K \overline{a} \overline{c} \overline{x} \rightarrow e$	alternative in pattern match
ψ	::=		value types
		D	data type (<i>not newtypes</i>)
		(\rightarrow)	arrow
		(\Rightarrow)	prop. arrow
		(\sim_ρ^κ)	equality
		$\forall a:\kappa.\tau$	polymorphism
		$\psi \tau$	application

Note that the value form $\Lambda a:\kappa.v$ requires that the expression inside the Λ itself be a value. This is because the operational semantics ignores type abstractions and evaluates under them.

A.2 Typing judgements

Note that the statement, for example, $a \# \Gamma$ means that the variable a is fresh in the context Γ .

$\boxed{\vdash \Gamma}$ Context validity

$$\frac{}{\vdash \emptyset} \text{CTX_EMPTY}$$

$$\frac{\vdash \Gamma \quad a \# \Gamma}{\vdash \Gamma, a : \kappa} \text{CTX_TYVAR}$$

$$\frac{\Gamma \vdash \tau \sim_{\rho} \sigma : \star \quad c \# \Gamma}{\vdash \Gamma, c : \phi} \text{CTX_COVAR}$$

$$\frac{\Gamma \vdash \tau : \star \quad x \# \Gamma}{\vdash \Gamma, x : \tau} \text{CTX_VAR}$$

 $\boxed{\Gamma \vdash \tau : \kappa}$ Type kinding

$$\frac{\vdash \Gamma \quad a : \kappa \in \Gamma}{\Gamma \vdash a : \kappa} \text{TY_VAR}$$

$$\frac{\Gamma \vdash \tau_1 : \kappa_1 \rightarrow \kappa_2 \quad \Gamma \vdash \tau_2 : \kappa_1}{\Gamma \vdash \tau_1 \tau_2 : \kappa_2} \text{TY_APP}$$

$$\frac{\vdash \Gamma \quad T : \kappa}{\Gamma \vdash T : \kappa} \text{TY_ADT}$$

$$\frac{\vdash \Gamma}{\Gamma \vdash (\rightarrow) : \star \rightarrow \star \rightarrow \star} \text{TY_ARROW}$$

$$\frac{\vdash \Gamma}{\Gamma \vdash (\Rightarrow) : \star \rightarrow \star \rightarrow \star} \text{TY_PROPARROW}$$

$$\frac{\vdash \Gamma}{\Gamma \vdash (\sim_{\rho}^{\kappa}) : \kappa \rightarrow \kappa \rightarrow \star} \text{TY_EQUALITY}$$

$$\frac{\Gamma, a : \kappa \vdash \tau : \star}{\Gamma \vdash \forall a : \kappa. \tau : \star} \text{TY_FORALL}$$

$$\frac{\vdash \Gamma \quad F : [\bar{a} : \kappa']. \kappa \quad \overline{\Gamma \vdash \tau : \kappa'}}{\Gamma \vdash F(\bar{\tau}) : \kappa} \text{TY_TYFUN}$$

 $\boxed{\Gamma \vdash e : \tau}$ Expression typing

$$\frac{\vdash \Gamma \quad x : \tau \in \Gamma}{\Gamma \vdash x : \tau} \text{TM_VAR}$$

$$\frac{\Gamma, x : \tau \vdash e : \sigma}{\Gamma \vdash \lambda x : \tau. e : \tau \rightarrow \sigma} \text{TM_ABS}$$

$$\frac{\Gamma \vdash e_1 : \tau \rightarrow \sigma \quad \Gamma \vdash e_2 : \tau}{\Gamma \vdash e_1 e_2 : \sigma} \text{TM_APP}$$

$$\begin{array}{c}
\frac{\Gamma, a:\kappa \vdash e : \tau}{\Gamma \vdash \Lambda a:\kappa.e : \forall a:\kappa.\tau} \quad \text{TM_TABS} \\
\frac{\Gamma \vdash e : \forall a:\kappa.\sigma}{\Gamma \vdash \tau : \kappa} \quad \text{TM_TAPP} \\
\frac{\Gamma, c:\sigma_1 \sim_\rho \sigma_2 \vdash e : \tau}{\Gamma \vdash \lambda c:\sigma_1 \sim_\rho \sigma_2.e : \phi \Rightarrow \tau} \quad \text{TM_CABS} \\
\frac{\Gamma \vdash e : (\sigma_1 \sim_\rho \sigma_2) \Rightarrow \tau \quad \Gamma \vdash \gamma : \sigma_1 \sim_\rho \sigma_2}{\Gamma \vdash e\gamma : \tau} \quad \text{TM_CAPP} \\
\frac{\vdash \Gamma \quad K : \tau}{\Gamma \vdash K : \tau} \quad \text{TM_DATA CON} \\
\Gamma \vdash e : D \bar{\sigma} \\
\Gamma \vdash \tau : \star \\
\forall alt_i \text{ s.t. } alt_i \in \overline{alt}, \\
alt_i = K_i \bar{a}_i \bar{c}_i \bar{x}_i \rightarrow e_i \\
K_i : \forall a'_i : \kappa_i. \forall b'_i : \kappa'_i. \bar{\phi}_i \Rightarrow \bar{\tau}_i \rightarrow D \bar{a}'_i \\
\frac{\Gamma, a_i : \kappa'_i, (\bar{c}_i : \bar{\phi}_i, \bar{x}_i : \bar{\tau}_i) [\sigma / a'_i] [a_i / b'_i] \vdash e_i : \tau}{alt \text{ is exhaustive}} \quad \text{TM_CASE} \\
\Gamma \vdash \mathbf{case}_\tau e \text{ of } \overline{alt} : \tau \\
\frac{\Gamma \vdash e : \tau_1 \quad \Gamma \vdash \gamma : \tau_1 \sim_R \tau_2}{\Gamma \vdash e \triangleright \gamma : \tau_2} \quad \text{TM_CAST} \\
\frac{\emptyset \vdash \gamma : H_1 \sim_N H_2 \quad H_1 \neq H_2 \quad \Gamma \vdash \tau : \star}{\Gamma \vdash \mathbf{contra} \gamma \tau : \tau} \quad \text{TM_CONTRA}
\end{array}$$

A.3 Small-step operational semantics

Because we evaluate under Λ , it is necessary to index the step relation by a typing environment. In practice, this environment will hold only type variables, never coercion or term variables.

$\boxed{e_1 \xrightarrow{\Gamma} e_2}$ Small-step operational semantics

$$\frac{}{(\lambda x:\tau.e_1) e_2 \xrightarrow{\Gamma} e_1[e_2/x]} \quad \text{S_BETA}$$

$$\frac{}{(\Lambda a:\kappa.v) \tau \xrightarrow{\Gamma} v[\tau/a]} \quad \text{S_TBETA}$$

$$\frac{}{(\lambda c:\phi.e) \gamma \xrightarrow{\Gamma} e[\gamma/c]} \quad \text{S_CBETA}$$

$$\frac{alt_i = K\bar{a}\bar{c}\bar{x} \rightarrow e'}{\mathbf{case}_{\tau_0} K\bar{\tau}\bar{\sigma}\bar{\gamma}\bar{e} \mathbf{of} \bar{alt} \xrightarrow{\Gamma} e'[\sigma/a][\gamma/c][e/x]} \quad \mathbf{S_IOTA}$$

$$\frac{}{(v \triangleright \gamma_1) \triangleright \gamma_2 \xrightarrow{\Gamma} v \triangleright (\gamma_1 \mathbin{\text{\textcircled{;}}} \gamma_2)} \quad \mathbf{S_TRANS}$$

$$\frac{e \xrightarrow{\Gamma} e'}{\Lambda a:\kappa.e \xrightarrow{\Gamma} \Lambda a:\kappa.e'} \quad \mathbf{S_TABS_CONG}$$

$$\frac{e_1 \xrightarrow{\Gamma} e'_1}{e_1 e_2 \xrightarrow{\Gamma} e'_1 e_2} \quad \mathbf{S_APP_CONG}$$

$$\frac{e \xrightarrow{\Gamma} e'}{e\tau \xrightarrow{\Gamma} e'\tau} \quad \mathbf{S_TAPP_CONG}$$

$$\frac{e \xrightarrow{\Gamma} e'}{e\gamma \xrightarrow{\Gamma} e'\gamma} \quad \mathbf{S_CAPP_CONG}$$

$$\frac{e \xrightarrow{\Gamma} e'}{\mathbf{case}_{\tau} e \mathbf{of} \bar{alt} \xrightarrow{\Gamma} \mathbf{case}_{\tau} e' \mathbf{of} \bar{alt}} \quad \mathbf{S_CASE_CONG}$$

$$\frac{e \xrightarrow{\Gamma} e'}{e \triangleright \gamma \xrightarrow{\Gamma} e' \triangleright \gamma} \quad \mathbf{S_CAST_CONG}$$

$$\frac{\eta_1 = \mathbf{sym}(\mathbf{nth}^1 \eta_0) \quad \eta_2 = \mathbf{nth}^2 \eta_0 \quad \Gamma \vdash v : \sigma_1 \rightarrow \sigma_2}{(v \triangleright \eta_0) e' \xrightarrow{\Gamma} v(e' \triangleright \eta_1) \triangleright \eta_2} \quad \mathbf{S_PUSH}$$

$$\frac{\Gamma \vdash v : \forall a:\kappa.\sigma' \quad \Gamma \vdash \tau : \kappa}{(v \triangleright \gamma)\tau \xrightarrow{\Gamma} v\tau \triangleright \gamma@ \tau} \quad \mathbf{S_TPUSH}$$

$$\frac{\eta_{11} = \mathbf{nth}^1(\mathbf{nth}^1 \eta_0) \quad \eta_{12} = \mathbf{nth}^2(\mathbf{nth}^1 \eta_0) \quad \eta_2 = \mathbf{nth}^2 \eta_0 \quad \gamma'' = \eta_{11} \mathbin{\text{\textcircled{;}}} \gamma' \mathbin{\text{\textcircled{;}}} \mathbf{sym} \eta_{12} \quad \Gamma \vdash v : (\sigma_1 \sim_{\rho}^{\kappa} \sigma_2) \Rightarrow \sigma_3 \quad \Gamma \vdash \gamma' : \sigma_4 \sim_{\rho}^{\kappa} \sigma_5}{(v \triangleright \eta_0) \gamma' \xrightarrow{\Gamma} v \gamma'' \triangleright \eta_2} \quad \mathbf{S_CPUSH}$$

$$\frac{}{\Lambda a:\kappa.(v \triangleright \gamma) \xrightarrow{\Gamma} (\Lambda a:\kappa.v) \triangleright (\forall a:\kappa.\gamma)} \quad \mathbf{S_APUSH}$$

$$\begin{array}{c}
\Gamma \vdash \eta : D \bar{\tau} \sim_{\mathbb{R}} D \bar{\tau}' \\
\hline
K : \forall \bar{a}:\bar{\kappa}.\forall \bar{b}:\bar{\kappa}'.(\sigma' \sim_{\rho} \sigma'') \Rightarrow \bar{\tau}'' \rightarrow D \bar{a} \\
\hline
\Gamma \vdash \gamma : (\sigma' \sim_{\rho} \sigma'')[\bar{\tau}/\bar{a}][\bar{\sigma}/\bar{b}] \\
\hline
\gamma' = \mathbf{sym}(\sigma'[\mathbf{nth} \eta/a]_{\rho}) \ddagger \gamma \ddagger \sigma''[\mathbf{nth} \eta/a]_{\rho} \\
\hline
e' = e \triangleright \tau''[\mathbf{nth} \eta/a]_{\mathbb{R}} \\
\hline
\mathbf{case}_{\tau_0}(K \bar{\tau} \bar{\sigma} \bar{\gamma} \bar{e}) \triangleright \eta \mathbf{of} \overline{\mathbf{alt}} \xrightarrow{\Gamma} \mathbf{case}_{\tau_0} K \bar{\tau}' \bar{\sigma}' \bar{\gamma}' \bar{e}' \mathbf{of} \overline{\mathbf{alt}} \quad \mathbf{S_KPUSH}
\end{array}$$

The typing context Γ threaded through this relation is used only in the premises of the “push” rules. Outside of $\mathbf{S_KPUSH}$, these premises are needed only to allow us to prove the preservation theorem (Theorem 34) without depending on consistency. In the $\mathbf{S_KPUSH}$ rule, however, the typing judgements are necessary to extract information used in the reduction; specifically, we need the $\bar{\tau}'$ from the type η , which appear in the reduct.

B Global context well-formedness

We assume throughout the paper and this appendix that the global context is well formed. Here, we explain precisely what can appear in the global context and what restrictions there are:

1. The global context may contain $C : [\bar{a}:\bar{\kappa}].\tau \sim_{\rho} \sigma$:
 - (a) $\overline{a:\kappa \vdash \tau : \kappa_0}$
 - (b) $\overline{a:\kappa \vdash \sigma : \kappa_0}$
2. The global context may contain $T : \kappa$.
3. The global context may contain $K : \tau$:
 - (a) $\tau = \forall \bar{a}:\bar{\kappa}.\forall \bar{b}:\bar{\kappa}'.\bar{\phi} \Rightarrow \bar{\sigma} \rightarrow D \bar{a}$
 - (b) $\emptyset \vdash \tau : \star$
4. The global context may contain $F : [\bar{a}:\bar{\kappa}].\kappa_0$.
5. For all H , $\mathit{roles}(H) \models H$.

C Properties of roles

Lemma 2 (Permutation of role checking). *If $\Omega \vdash \tau : \rho$ and Ω' is a permutation of Ω , then $\Omega' \vdash \tau : \rho$.*

Lemma 3 (Weakening of role checking). *If $\Omega \vdash \tau : \rho$, then $\Omega, a:\rho' \vdash \tau : \rho$.*

Lemma 4 (Strengthening of role checking). *If $\Omega, a:\rho' \vdash \tau : \rho$ and a does not appear free in τ , then $\Omega \vdash \tau : \rho$.*

Lemma 5 (Nominal roles are infectious). *Let \bar{a} be the free variables in σ . We have $\Omega \vdash \sigma : \mathbb{N}$ if and only if every $a_i \in \bar{a}$ is at role \mathbb{N} in Ω .*

Lemma 6 (Sub-roling). *If $\Omega \vdash \tau : \rho$ and $\rho \leq \rho'$, then $\Omega \vdash \tau : \rho'$.*

This next property is not needed for type safety, but it says that it is always sound to assign stricter roles to the variables in a type declaration.

Lemma 7 (Context sub-roling). *If $\overline{a}:\overline{\rho} \vdash \tau : \rho_0$ and $\overline{\rho'} \leq \overline{\rho}$, then $\overline{a}:\overline{\rho'} \vdash \tau : \rho_0$.*

Lemma 8 (Roles on type constants). *For any Ω, H , and ρ , $\Omega \vdash H : \rho$.*

Proof. By case analysis on ρ :

Case $\rho = P$: By RTY_PHANTOM.

Case $\rho = R$: By RTY_TYCONAPP.

Case $\rho = N$: By RTY_TYCON.

□

D Structural properties

D.1 Weakening

Let bnd be a metavariable for a context binding. That is:

$$\begin{array}{l} bnd ::= a:\kappa \\ \quad | c:\phi \\ \quad | x:\tau \end{array}$$

Lemma 9 (Type kinding weakening). *If $\Gamma, \Gamma' \vdash \tau : \kappa$ and $\vdash \Gamma, bnd, \Gamma'$, then $\Gamma, bnd, \Gamma' \vdash \tau : \kappa$.*

Lemma 10 (Coercion typing weakening). *If $\Gamma, \Gamma' \vdash \gamma : \phi$ and $\vdash \Gamma, bnd, \Gamma'$, then $\Gamma, bnd, \Gamma' \vdash \gamma : \phi$.*

Lemma 11 (Term typing weakening). *If $\Gamma, \Gamma' \vdash e : \tau$ and $\vdash \Gamma, bnd, \Gamma'$, then $\Gamma, bnd, \Gamma' \vdash e : \tau$.*

D.2 Substitution

Lemma 12 (Type variable substitution). *Suppose $\Gamma \vdash \sigma : \kappa_1$. Then:*

1. *If $\vdash \Gamma, a:\kappa_1, \Gamma'$, then $\vdash \Gamma, \Gamma'[\sigma/a]$;*
2. *If $\Gamma, a:\kappa_1, \Gamma' \vdash \tau : \kappa_2$, then $\Gamma, \Gamma'[\sigma/a] \vdash \tau[\sigma/a] : \kappa_2$.*

Lemma 13 (Type variable substitution in coercions). *If $\Gamma, a:\kappa, \Gamma' \vdash \gamma : \phi$ and $\Gamma \vdash \sigma : \kappa$, then $\Gamma, \Gamma'[\sigma/a] \vdash \gamma[\sigma/a] : \phi[\sigma/a]$.*

Lemma 14 (Type variable substitution in terms). *If $\Gamma, a:\kappa, \Gamma' \vdash e : \tau$ and $\Gamma \vdash \sigma : \kappa$, then $\Gamma, \Gamma'[\sigma/a] \vdash e[\sigma/a] : \tau[\sigma/a]$.*

Lemma 15 (Coercion strengthening).

1. *If $\vdash \Gamma, c:\phi, \Gamma'$, then $\vdash \Gamma, \Gamma'$;*
2. *If $\Gamma, c:\phi, \Gamma' \vdash \tau : \kappa$, then $\Gamma, \Gamma' \vdash \tau : \kappa$.*

Lemma 16 (Coercion substitution). *If $\Gamma, c:\phi_1, \Gamma' \vdash \gamma : \phi_2$ and $\Gamma \vdash \eta : \phi_1$, then $\Gamma, \Gamma' \vdash \gamma[\eta/c] : \phi_2$.*

Lemma 17 (Coercion substitution in terms). *If $\Gamma, c:\phi, \Gamma' \vdash e : \tau$ and $\Gamma \vdash \eta : \phi$, then $\Gamma, \Gamma' \vdash e[\eta/c] : \tau$.*

Lemma 18 (Term strengthening).

1. If $\vdash \Gamma, x:\tau, \Gamma'$, then $\vdash \Gamma, \Gamma'$;
2. If $\Gamma, x:\tau, \Gamma' \vdash \sigma : \kappa$, then $\Gamma, \Gamma' \vdash \sigma : \kappa$.

Lemma 19 (Term strengthening in coercions). If $\Gamma, x:\tau, \Gamma' \vdash \gamma : \phi$, then $\Gamma, \Gamma' \vdash \gamma : \phi$.

Lemma 20 (Term substitution). If $\Gamma, x:\sigma, \Gamma' \vdash e : \tau$ and $\Gamma \vdash e' : \sigma$, then $\Gamma, \Gamma' \vdash e[e'/x] : \tau$.

D.3 Context regularity

Lemma 21 (Type context regularity for types). If $\Gamma \vdash \tau : \kappa$, then $\vdash \Gamma$.

Lemma 22 (Coercion context regularity). If $\Gamma \vdash \gamma : \phi$, then $\vdash \Gamma$.

Lemma 23 (Term context regularity). If $\Gamma \vdash e : \tau$, then $\vdash \Gamma$.

D.4 Classifier regularity

Lemma 24 (Coercion typing regularity). If $\Gamma \vdash \gamma : \tau \sim_\rho \sigma$, then $\Gamma \vdash \tau \sim_\rho \sigma : \star$.

Lemma 25 (Coercion homogeneity). If $\Gamma \vdash \gamma : \tau \sim_\rho \sigma$, then $\Gamma \vdash \tau : \kappa$ and $\Gamma \vdash \sigma : \kappa$.

Proof. Direct from Lemma 24. □

Lemma 26 (Term typing regularity). If $\Gamma \vdash e : \tau$, then $\Gamma \vdash \tau : \star$.

D.5 Determinacy

Lemma 27 (Uniqueness of type kinding). If $\Gamma \vdash \tau : \kappa_1$ and $\Gamma \vdash \tau : \kappa_2$ then $\kappa_1 = \kappa_2$.

Lemma 28 (Uniqueness of coercion typing). If $\Gamma \vdash c : \phi_1$ and $\Gamma \vdash c : \phi$ then $\phi_1 = \phi$.

Lemma 29 (Uniqueness of term typing). If $\Gamma \vdash e : \tau_1$ and $\Gamma \vdash e : \tau_2$ then $\tau_1 = \tau_2$.

Lemma 30 (Values do not step). For all v , there exists no e such that $v \xrightarrow{\Gamma} e$.

Lemma 31 (Coerced values do not step). For all v and γ , there exists no e such that $(v \triangleright \gamma) \xrightarrow{\Gamma} e$.

Lemma 32 (Determinacy of evaluation). If $e \xrightarrow{\Gamma} e_1$ and $e \xrightarrow{\Gamma} e_2$ then $e_1 = e_2$.

E Preservation

E.1 Lifting

Lifting is a process that transports coercions through a type. This operation is used by the rule `S_KPUSH` to push a coercion into the arguments of a data constructor. It is defined by the following algorithm, with patterns to be tried in order from top to bottom. Note that the context Γ is an implicit argument of this function.

$$\begin{array}{lll}
 \tau[\overline{\gamma/b}]_{\mathsf{P}} & = \langle \tau[\overline{\sigma/b}], \tau[\overline{\sigma'/b}] \rangle_{\mathsf{P}} & (\Gamma \vdash \gamma : \sigma \sim_{\rho} \sigma') \\
 a[\overline{\gamma/b}]_{\rho} & = \gamma_i & (a = b_i \wedge \Gamma \vdash \gamma_i : \sigma \sim_{\rho} \sigma') \\
 a[\overline{\gamma/b}]_{\mathsf{R}} & = \mathbf{sub} \gamma_i & (a = b_i) \\
 a[\overline{\gamma/b}]_{\mathsf{N}} & = \langle a \rangle & (a \notin \overline{b}) \\
 a[\overline{\gamma/b}]_{\mathsf{R}} & = \mathbf{sub} \langle a \rangle & (a \notin \overline{b}) \\
 (H \overline{\tau})[\overline{\gamma/b}]_{\mathsf{R}} & = H(\overline{\tau[\overline{\gamma/b}]_{\rho}}) & (\overline{\rho} \text{ is a prefix of } \mathit{roles}(H)) \\
 H[\overline{\gamma/b}]_{\mathsf{N}} & = \langle H \rangle & \\
 (\tau_1 \tau_2)[\overline{\gamma/b}]_{\rho} & = \tau_1[\overline{\gamma/b}]_{\rho} \tau_2[\overline{\gamma/b}]_{\mathsf{N}} & \\
 (\forall a:\kappa.\tau)[\overline{\gamma/b}]_{\rho} & = \forall a:\kappa.\tau[\overline{\gamma/b}]_{\rho} & \\
 (F(\overline{\tau}))[\overline{\gamma/b}]_{\mathsf{N}} & = F(\overline{\tau[\overline{\gamma/b}]_{\mathsf{N}}}) & \\
 (F(\overline{\tau}))[\overline{\gamma/b}]_{\mathsf{R}} & = \mathbf{sub} F(\overline{\tau[\overline{\gamma/b}]_{\mathsf{N}}}) &
 \end{array}$$

Lemma 33 (Lifting). *If:*

1. $\Gamma \vdash \gamma : H \overline{\tau} \sim_{\mathsf{R}} H \overline{\sigma}$;
2. $\overline{\Gamma} \vdash \tau : \kappa$;
3. $\overline{\Gamma} \vdash \sigma : \kappa$;
4. H is not a **newtype**;
5. $\Omega \vdash \sigma_0 : \rho_0$, where $\overline{b'}$ is the type variables in Γ, Γ'
 $\Omega = \overline{b':\mathsf{N}}, \overline{b} : \mathit{roles}(H)$;
6. $\Gamma, \overline{b}:\kappa, \Gamma' \vdash \sigma_0 : \kappa'$; and
7. Γ' contains only type variable bindings.

then:

$$\Gamma, \Gamma' \vdash \sigma_0[\mathbf{nth} \overline{\gamma/b}]_{\rho_0} : \sigma_0[\overline{\tau/b}] \sim_{\rho_0} \sigma_0[\overline{\sigma/b}]$$

Proof. First, because Γ' contains only type variable bindings, then a type variable substitution has no effect on Γ' (which can contain only *kinds*).

If $\rho_0 = \mathsf{P}$, then the first equation of the algorithm matches, and we have $\sigma_0[\mathbf{nth} \overline{\gamma/b}]_{\mathsf{P}} = \langle \sigma_0[\overline{\tau/b}], \sigma_0[\overline{\sigma/b}] \rangle_{\mathsf{P}}$, and we are done, applying Lemma 12.

So, we assume now that $\rho_0 \neq \mathsf{P}$.

Let $\overline{\rho} = \mathit{roles}(H)$.

We proceed by induction on the derivation of $\Gamma, \overline{b}:\kappa, \Gamma' \vdash \sigma_0 : \kappa'$. Each case concludes by the application of the appropriate substitution lemma(s).

Case TY_VAR: We know $\sigma_0 = a$.

Case ($a = b_i$):

Case ($\rho_0 = \rho_i$): Here, we have $\sigma_0[\mathbf{nth} \overline{\gamma/b}]_{\rho_0} = \mathbf{nth}^i \gamma$, $\sigma_0[\overline{\tau/b}] = \tau_i$, and $\sigma_0[\overline{\sigma/b}] = \sigma_i$. Thus, we are done, by CO_NTH.

Case ($\rho_0 = \mathsf{R}, \rho_i = \mathsf{N}$): Similar to the last case, fixing the roles with a use of **sub**.

Case ($\rho_0 = \mathsf{N}, \rho_i \neq \mathsf{N}$): This case is impossible. We know $\Omega \vdash a : \mathsf{N}$. By inversion then, we know $a:\mathsf{N} \in \Omega$. Yet, we know that ρ_i is the i th role in $\mathit{roles}(H)$, and by the definition of Ω , $a:\rho_i \in \Omega$. This contradicts $\rho_i \neq \mathsf{N}$, and we are done.

Case ($a \notin \bar{b}$):

Case ($\rho_0 = \mathbf{N}$): Here, $\sigma_0[\overline{\mathbf{nth}} \gamma / \bar{b}]_{\mathbf{N}} = \langle \sigma_0 \rangle$, $\sigma_0[\bar{\tau} / \bar{b}] = \sigma_0$, and $\sigma_0[\overline{\sigma} / \bar{b}] = \sigma_0$, so we are done, by `CO_REFL`.

Case ($\rho_0 = \mathbf{R}$): Similar to last case, fixing the output role with `sub`.

Case `TY_APP`:

Case ($\sigma_0 = H' \bar{\sigma}'$, $\rho_0 = \mathbf{R}$): Here $(H' \bar{\sigma}')[\overline{\mathbf{nth}} \gamma / \bar{b}]_{\mathbf{R}} = H'(\overline{\sigma'[\mathbf{nth} \gamma / \bar{b}]_{\rho'}}$), where $\bar{\rho}'$ is a prefix of $\text{roles}(H')$. Let $\eta = H'(\overline{\sigma'[\mathbf{nth} \gamma / \bar{b}]_{\rho'}}$. Then, we must show $\Gamma, \Gamma' \vdash \eta : H' \bar{\sigma}'[\bar{\tau} / \bar{b}] \sim_{\mathbf{R}} H' \bar{\sigma}'[\overline{\sigma} / \bar{b}]$. We will use `CO_TYCONAPP`. We must show

$$\Gamma, \Gamma' \vdash \overline{\sigma'[\mathbf{nth} \gamma / \bar{b}]_{\rho'}} : \sigma'[\bar{\tau} / \bar{b}] \sim_{\rho'} \sigma'[\overline{\sigma} / \bar{b}].$$

We do this by induction, for each $\sigma'_i \in \bar{\sigma}'$. All of the premises of the lifting lemma are satisfied automatically, except for premise 5. Fix i . We must show $\Omega \vdash \sigma'_i : \rho'_i$. We know $\Omega \vdash H' \bar{\sigma}' : \mathbf{R}$. This can be proved by either `RTY_TYCONAPP` or `RTY_APP`. If it is by the former, we are done by inversion. If it is by the latter, then we know $\Omega \vdash \sigma'_i : \mathbf{N}$. We apply Lemma 6, and we are done.

Other applications: Apply the induction hypothesis. Premise 5 of the lifting lemma is satisfied by correspondence between `RTY_APP` and `CO_APP`.

Case `TY_ADT`:

Case ($\rho_0 = \mathbf{N}$): Here $H[\overline{\mathbf{nth}} \gamma / \bar{b}]_{\mathbf{N}} = \langle H \rangle$, and we are done by `CO_REFL`.

Case ($\rho_0 = \mathbf{R}$): Here $H[\overline{\mathbf{nth}} \gamma / \bar{b}]_{\mathbf{R}} = H(\emptyset)$ and we are done by `CO_TYCONAPP`.

Cases `TY_ARROW`, `TY_EQUALITY`: Similar to `TY_ADT`.

Case `TY_FORALL`: By the induction hypothesis. Note that the roles in `RTY_FORALL` and `CO_FORALL` match up, and that the new binding in `RTY_FORALL` is given a nominal role, echoed in the definition of Ω in this lemma's premises.

Case `TY_TYFUN`: By the induction hypothesis, once again noting the correspondence between `RTY_TYFAM` and `CO_TYFAM`.

□

E.2 Preservation

Theorem 34 (Preservation). *If $\Gamma \vdash e : \tau$ and $e \xrightarrow{\Gamma} e'$, then $\Gamma \vdash e' : \tau$.*

Proof. By induction on the derivation of $e \xrightarrow{\Gamma} e'$.

Beta rules: By substitution.

Case `S_IOTA`: We know $\Gamma \vdash \text{case}_{\tau_0} K \bar{\tau} \bar{\sigma} \bar{\gamma} \bar{e}$ of $\bar{alt} : \tau_0$, where $\text{alt}_i = K \bar{a} \bar{c} \bar{x} \rightarrow e'$. We must show $\Gamma \vdash e'[\overline{\sigma} / \bar{a}][\bar{\gamma} / \bar{c}][\bar{e} / \bar{x}] : \tau_0$. By inversion on `TM_CASE`, we see

$$\begin{aligned} & \Gamma \vdash K \bar{\tau} \bar{\sigma} \bar{\gamma} \bar{e} : D \bar{\tau} \\ & K : \forall \bar{a}' : \kappa. \forall \bar{b}' : \kappa'. \bar{\phi} \Rightarrow \bar{\tau}' \rightarrow D \bar{a}' \\ & \Gamma, \bar{a} : \kappa', \bar{c} : \phi[\bar{\tau} / \bar{a}][\bar{a} / \bar{b}'], \bar{x} : \tau'[\bar{\tau} / \bar{a}][\bar{a} / \bar{b}'] \vdash e' : \tau_0 \end{aligned}$$

We also know that $\Gamma \vdash \tau_0 : \star$, which implies that none of the variables \bar{a} are mentioned in τ_0 . We can do induction on the length of $\bar{\tau}$ to see that

$$\Gamma \vdash K\bar{\tau} : \forall \bar{b}' : \kappa'. \bar{\phi}[\bar{\tau}/\bar{a}'] \Rightarrow \bar{\tau}'[\bar{\tau}/\bar{a}'] \rightarrow D \bar{a}'[\bar{\tau}/\bar{a}']$$

This simplifies to

$$\Gamma \vdash K\bar{\tau} : \forall \bar{b}' : \kappa'. \bar{\phi}[\bar{\tau}/\bar{a}'] \Rightarrow \bar{\tau}'[\bar{\tau}/\bar{a}'] \rightarrow D \bar{\tau}$$

Now, we do induction on the length of $\bar{\sigma}$ to see that

$$\Gamma \vdash K\bar{\tau}\bar{\sigma} : \bar{\phi}[\bar{\tau}/\bar{a}'][\bar{\sigma}/\bar{b}'] \Rightarrow \bar{\tau}'[\bar{\tau}/\bar{a}'][\bar{\sigma}/\bar{b}'] \rightarrow D \bar{\tau}$$

and

$$\overline{\Gamma \vdash \sigma : \kappa'}$$

We can then use repeated application of the type variable substitution lemma to get

$$\overline{\Gamma, c : \bar{\phi}[\bar{\tau}/\bar{a}'][\bar{\sigma}/\bar{b}'], x : \bar{\tau}'[\bar{\tau}/\bar{a}'][\bar{\sigma}/\bar{b}'] \vdash e'[\bar{\sigma}/\bar{a}] : \tau_0}$$

using the following facts

$$\begin{aligned} \tau_0[\bar{\sigma}/\bar{a}] &= \tau_0 \\ \bar{\phi}[\bar{\tau}/\bar{a}'][\bar{a}/\bar{b}'][\bar{\sigma}/\bar{a}] &= \bar{\phi}[\bar{\tau}/\bar{a}'][\bar{\sigma}/\bar{b}'] \\ \bar{\tau}'[\bar{\tau}/\bar{a}'][\bar{a}/\bar{b}'][\bar{\sigma}/\bar{a}] &= \bar{\tau}'[\bar{\tau}/\bar{a}'][\bar{\sigma}/\bar{b}'] \end{aligned}$$

So, we have

$$\overline{\Gamma, c : \bar{\phi}[\bar{\tau}/\bar{a}'][\bar{\sigma}/\bar{b}'], x : \bar{\tau}'[\bar{\tau}/\bar{a}'][\bar{\sigma}/\bar{b}'] \vdash e'[\bar{\sigma}/\bar{a}] : \tau_0}$$

Starting from the type of $K\bar{\tau}\bar{\sigma}$, we do induction on the length of $\bar{\gamma}$ to get

$$\Gamma \vdash K\bar{\tau}\bar{\sigma}\bar{\gamma} : \bar{\tau}'[\bar{\tau}/\bar{a}'][\bar{\sigma}/\bar{b}'] \rightarrow D \bar{\tau}$$

and

$$\overline{\Gamma \vdash \gamma : \bar{\phi}[\bar{\tau}/\bar{a}'][\bar{\sigma}/\bar{b}']}$$

Thus, we can use the coercion variable substitution lemma to get

$$\overline{\Gamma, x : \bar{\tau}'[\bar{\tau}/\bar{a}'][\bar{\sigma}/\bar{b}'] \vdash e'[\bar{\sigma}/\bar{a}][\bar{\gamma}/\bar{c}] : \tau_0}$$

Finally we use analogous reasoning for term arguments \bar{e} to conclude

$$\Gamma \vdash e'[\bar{\sigma}/\bar{a}][\bar{\gamma}/\bar{c}][\bar{e}/\bar{x}] : \tau_0$$

as desired.

Case S_TRANS: We know that $\Gamma \vdash (v \triangleright \gamma_1) \triangleright \gamma_2 : \tau$ and need to show that $\Gamma \vdash v \triangleright (\gamma_1 \circ \gamma_2) : \tau$. Inversion gives us $\Gamma \vdash v : \sigma_1$, $\Gamma \vdash \gamma_1 : \sigma_1 \sim_R \sigma_2$, and $\Gamma \vdash \gamma_2 : \sigma_2 \sim_R \tau$. Straightforward use of typing rules shows that $\Gamma \vdash v \triangleright (\gamma_1 \circ \gamma_2) : \tau$, as desired.

Congruence rules: By induction.

Case S_PUSH: We adopt the variable names from the statement of the rule:

$$\frac{\eta_1 = \mathbf{sym}(\mathbf{nth}^1 \eta_0) \quad \eta_2 = \mathbf{nth}^2 \eta_0 \quad \Gamma \vdash v : \sigma_1 \rightarrow \sigma_2}{(v \triangleright \eta_0) e' \xrightarrow[\Gamma]{} v(e' \triangleright \eta_1) \triangleright \eta_2} \text{ S_PUSH}$$

We know that $\Gamma \vdash (v \triangleright \eta_0) e' : \sigma_4$ and we must show $\Gamma \vdash (v(e' \triangleright \eta_1)) \triangleright \eta_2 : \sigma_4$. Inversion tells us that $\Gamma \vdash \eta_0 : (\sigma_1 \rightarrow \sigma_2) \sim_{\mathbf{R}} (\sigma_3 \rightarrow \sigma_4)$ and $\Gamma \vdash e' : \sigma_3$. We can now see that $\Gamma \vdash \eta_1 : \sigma_3 \sim_{\mathbf{R}} \sigma_1$ and $\Gamma \vdash \eta_2 : \sigma_2 \sim_{\mathbf{R}} \sigma_4$. Thus, $\Gamma \vdash e' \triangleright \eta_1 : \sigma_1$ and $\Gamma \vdash v(e' \triangleright \eta_1) \triangleright \eta_2 : \sigma_4$ as desired.

Case S_TPUSH: We adopt the variable names from the statement of the rule:

$$\frac{\Gamma \vdash v : \forall a : \kappa . \sigma' \quad \Gamma \vdash \tau : \kappa}{(v \triangleright \gamma) \tau \xrightarrow[\Gamma]{} v \tau \triangleright \gamma @ \tau} \text{ S_TPUSH}$$

We know that $\Gamma \vdash (v \triangleright \gamma) \tau : \tau'$ and we must show that $\Gamma \vdash v \tau \triangleright \gamma @ \tau : \tau'$. Inversion tells us that $\Gamma \vdash \gamma : (\forall a : \kappa . \sigma') \sim_{\mathbf{R}} (\forall a : \kappa . \sigma'')$ where $\tau' = \sigma''[\tau/a]$. We can see that $\Gamma \vdash \gamma @ \tau : \sigma'[\tau/a] \sim_{\mathbf{R}} \sigma''[\tau/a]$ and thus that $\Gamma \vdash v \tau \triangleright \gamma @ \tau : \tau'$ as desired.

Case S_CPUSH: We adopt the variables names from the statement of the rule:

$$\frac{\eta_{11} = \mathbf{nth}^1(\mathbf{nth}^1 \eta_0) \quad \eta_{12} = \mathbf{nth}^2(\mathbf{nth}^1 \eta_0) \quad \eta_2 = \mathbf{nth}^2 \eta_0 \quad \gamma'' = \eta_{11} \circ \gamma' \circ \mathbf{sym} \eta_{12} \quad \Gamma \vdash v : (\sigma_1 \sim_{\rho}^{\kappa} \sigma_2) \Rightarrow \sigma_3 \quad \Gamma \vdash \gamma' : \sigma_4 \sim_{\rho}^{\kappa} \sigma_5}{(v \triangleright \eta_0) \gamma' \xrightarrow[\Gamma]{} v \gamma'' \triangleright \eta_2} \text{ S_CPUSH}$$

We know that $\Gamma \vdash (v \triangleright \eta_0) \gamma' : \sigma_6$ and we must show that $\Gamma \vdash v \gamma'' \triangleright \eta_2 : \sigma_6$. Inversion tells us that $\Gamma \vdash \eta_0 : (\sigma_1 \sim_{\rho} \sigma_2 \Rightarrow \sigma_3) \sim_{\mathbf{R}} (\sigma_4 \sim_{\rho} \sigma_5 \Rightarrow \sigma_6)$. We can now see the following:

$$\begin{aligned} \Gamma \vdash \eta_{11} : \sigma_1 \sim_{\rho} \sigma_4 \\ \Gamma \vdash \eta_{12} : \sigma_2 \sim_{\rho} \sigma_5 \\ \Gamma \vdash \eta_2 : \sigma_3 \sim_{\rho} \sigma_6 \\ \Gamma \vdash \gamma'' : \sigma_1 \sim_{\rho} \sigma_2 \end{aligned}$$

Thus $\Gamma \vdash v \gamma'' \triangleright \eta_2 : \sigma_6$ as desired.

Case S_APUSH: We adopt the variable names from the statement of the rule:

$$\frac{}{\Lambda a : \kappa . (v \triangleright \gamma) \xrightarrow[\Gamma]{} (\Lambda a : \kappa . v) \triangleright (\forall a : \kappa . \gamma)} \text{ S_APUSH}$$

Inversion tells us that there must be σ and σ' such that $\Gamma, a : \kappa \vdash v : \sigma$ and $\Gamma, a : \kappa \vdash \gamma : \sigma \sim_{\mathbf{R}} \sigma'$. We can then see that $\Gamma \vdash \Lambda a : \kappa . v : \forall a : \kappa . \sigma$ and $\Gamma \vdash \forall a : \kappa . \gamma : (\forall a : \kappa . \sigma) \sim_{\mathbf{R}} (\forall a : \kappa . \sigma')$. We thus get $\Gamma \vdash (\Lambda a : \kappa . v) \triangleright (\forall a : \kappa . \gamma) : \forall a : \kappa . \sigma'$ as desired.

Case S_KPUSH: We adopt the variable names from the statement of S_KPUSH:

$$\frac{\frac{\frac{\frac{\Gamma \vdash \eta : D \bar{\tau} \sim_R D \bar{\tau}'}{K : \forall \bar{a} : \bar{\kappa}. \forall \bar{b} : \bar{\kappa}'. (\sigma' \sim_\rho \sigma'') \Rightarrow \bar{\tau}'' \rightarrow D \bar{a}}{\Gamma \vdash \gamma : (\sigma' \sim_\rho \sigma'') [\bar{\tau}/\bar{a}] [\bar{\sigma}/\bar{b}]}}{\gamma' = \mathbf{sym} (\sigma' [\mathbf{nth} \eta / \bar{a}]_\rho) \ddagger \gamma \ddagger \sigma'' [\mathbf{nth} \eta / \bar{a}]_\rho}}{e' = e \triangleright \tau'' [\mathbf{nth} \eta / \bar{a}]_R}}{\mathbf{case}_{\tau_0} (K \bar{\tau} \bar{\sigma} \bar{\gamma} \bar{e}) \triangleright \eta \mathbf{of} \mathbf{alt} \rightarrow \mathbf{case}_{\tau_0} K \bar{\tau}' \bar{\sigma} \bar{\gamma}' \bar{e}' \mathbf{of} \mathbf{alt}} \quad \mathbf{S_KPUSH}$$

Inversion gives us the premises of this rule. We also know $\Gamma \vdash (K \bar{\tau} \bar{\sigma} \bar{\gamma} \bar{e}) \triangleright \eta : D \bar{\tau}'$. We must show $\Gamma \vdash (K \bar{\tau}' \bar{\sigma} \bar{\gamma}' \bar{e}') : D \bar{\tau}'$. Note that τ_0 and the \mathbf{alt} do not change, so we need not worry about them here.

Let $\bar{\phi} = \overline{(\sigma' \sim_\rho \sigma'')}$. From repeated inversion (and induction on the length of $\bar{\tau}$), we can derive

$$\overline{\Gamma \vdash \tau : \kappa}$$

Then, from homogeneity of coercions (Lemma 25) (and more induction on $\bar{\tau}'$), we see that

$$\overline{\Gamma \vdash \tau' : \kappa}$$

Putting this together, we get

$$\Gamma \vdash K \bar{\tau}' : (\forall \bar{b} : \bar{\kappa}'. \bar{\phi} \Rightarrow \bar{\tau}'' \rightarrow D \bar{a}) [\bar{\tau}'/\bar{a}]$$

or

$$\Gamma \vdash K \bar{\tau}' : \forall \bar{b} : \bar{\kappa}'. \bar{\phi} [\bar{\tau}'/\bar{a}] \Rightarrow \bar{\tau}'' [\bar{\tau}'/\bar{a}] \rightarrow D \bar{\tau}'$$

Taking $K \bar{\tau} \bar{\sigma} \bar{\gamma} \bar{e}$ apart further (and induction on $\bar{\sigma}$) tells us

$$\overline{\Gamma \vdash \sigma : \kappa'}$$

and thus that

$$\Gamma \vdash K \bar{\tau}' \bar{\sigma} : \bar{\phi} [\bar{\tau}'/\bar{a}] [\bar{\sigma}/\bar{b}] \Rightarrow \bar{\tau}'' [\bar{\tau}'/\bar{a}] [\bar{\sigma}/\bar{b}] \rightarrow D \bar{\tau}' [\bar{\sigma}/\bar{b}]$$

But, from $\overline{\Gamma \vdash \tau' : \kappa}$, we see that \bar{b} do not appear in $\bar{\tau}'$. So, we have

$$\Gamma \vdash K \bar{\tau}' \bar{\sigma} : \bar{\phi} [\bar{\tau}'/\bar{a}] [\bar{\sigma}/\bar{b}] \Rightarrow \bar{\tau}'' [\bar{\tau}'/\bar{a}] [\bar{\sigma}/\bar{b}] \rightarrow D \bar{\tau}'$$

Using techniques similar to that for $\bar{\tau}$ and $\bar{\sigma}$, we can derive the following:

$$\frac{\overline{\Gamma \vdash \gamma : \phi [\bar{\tau}/\bar{a}] [\bar{\sigma}/\bar{b}]}}{\overline{\Gamma \vdash e : \tau'' [\bar{\tau}/\bar{a}] [\bar{\sigma}/\bar{b}]}}$$

We need to conclude the following:

$$\frac{\overline{\Gamma \vdash \gamma' : \phi [\bar{\tau}'/\bar{a}] [\bar{\sigma}/\bar{b}]}}{\overline{\Gamma \vdash e' : \tau'' [\bar{\tau}'/\bar{a}] [\bar{\sigma}/\bar{b}]}}$$

We wish to use the lifting lemma (Lemma 33) to get types for $\overline{\sigma' [\mathbf{nth} \eta / \bar{a}]_\rho}$ and $\overline{\sigma'' [\mathbf{nth} \eta / \bar{a}]_\rho}$. So, we must first establish the premises of the lifting lemma.

1. $\Gamma \vdash \eta : D \bar{\tau} \sim_{\mathbb{R}} D \bar{\tau}'$, from the inversion on S_KPUSH (and weakening to change the context);
2. $\bar{\Gamma} \vdash \tau : \kappa$, as above;
3. $\bar{\Gamma} \vdash \tau' : \kappa$, as above;
4. D is not a **newtype**: by choice of metavariable.
5. $\bar{\Omega} \vdash \sigma' : \rho$ and $\bar{\Omega} \vdash \sigma'' : \rho$: Here, $\bar{\Omega} = \bar{b}':\mathbb{N}, \bar{a} : \text{roles}(D)$ where \bar{b}' are the type variables bound in Γ , along with the existential variables \bar{b} . (That is, the Γ' in the statement of the lifting lemma is $\bar{b}:\kappa'$.) By ROLES_DATA, we can see that $\bar{\Omega} \vdash (\sigma' \sim_{\rho} \sigma'') : \mathbb{R}$. This can be established by either RTY_TYCONAPP or by RTY_APP. In the former case, we get the desired outcome by looking at ROLES_EQUALITY. In the latter case, we see that $\bar{\Omega} \vdash \sigma'_i : \mathbb{N}$ or $\bar{\Omega} \vdash \sigma''_i : \mathbb{N}$ and then use role subsumption (Lemma 6).
6. $\Gamma, \bar{a}:\bar{\kappa}, \bar{b}:\kappa' \vdash \sigma' : \kappa''$ and the same for σ'' : This comes from the well-formedness of the global context, including the type of K .
7. $\bar{b}:\kappa'$ must contain only type variable bindings: It sure does.

Now, we can conclude

$$\frac{\overline{\Gamma, b:\kappa' \vdash \sigma' [\mathbf{nth} \eta / \bar{a}]_{\rho} : \sigma' [\bar{\tau} / \bar{a}] \sim_{\rho} \sigma' [\bar{\tau}' / \bar{a}]}}{\overline{\Gamma, b:\kappa' \vdash \sigma'' [\mathbf{nth} \eta / \bar{a}]_{\rho} : \sigma'' [\bar{\tau} / \bar{a}] \sim_{\rho} \sigma'' [\bar{\tau}' / \bar{a}]}}$$

We then do type variable substitution to get

$$\frac{\overline{\Gamma \vdash \sigma' [\mathbf{nth} \eta / \bar{a}]_{\rho} [\sigma / \bar{b}] : \sigma' [\bar{\tau} / \bar{a}] [\sigma / \bar{b}] \sim_{\rho} \sigma' [\bar{\tau}' / \bar{a}] [\sigma / \bar{b}]}}{\overline{\Gamma \vdash \sigma'' [\mathbf{nth} \eta / \bar{a}]_{\rho} [\sigma / \bar{b}] : \sigma'' [\bar{\tau} / \bar{a}] [\sigma / \bar{b}] \sim_{\rho} \sigma'' [\bar{\tau}' / \bar{a}] [\sigma / \bar{b}]}}$$

Now, by CO_TRANS, we can conclude

$$\overline{\Gamma \vdash \gamma' : \phi [\bar{\tau}' / \bar{a}] [\sigma / \bar{b}]}$$

as desired.

To type the \bar{e}' , we need to apply the lifting lemma once again, this time to $\tau'' [\mathbf{nth} \eta / \bar{a}]_{\mathbb{R}}$. Much of our work at establishing premises carries over, except for these:

5. $\bar{\Omega} \vdash \tau'' : \mathbb{R}$ (with $\bar{\Omega}$ as above): This comes directly from the premises of ROLES_DATA, noting that $\bar{\tau}''$ appears in as an argument type to K .
6. $\Gamma, \bar{a}:\bar{\kappa}, \bar{b}:\kappa' \vdash \tau'' : \kappa''$: This comes from the well-formedness of the global context, including the type of K .

We then apply the lifting lemma to conclude that

$$\overline{\Gamma, \bar{b}:\kappa' \vdash \tau'' [\mathbf{nth} \gamma / \bar{a}]_{\mathbb{R}} : \tau'' [\bar{\tau} / \bar{a}] \sim_{\mathbb{R}} \tau'' [\bar{\tau}' / \bar{a}]}$$

We use type variable substitution to get

$$\overline{\Gamma \vdash \tau'' [\mathbf{nth} \gamma / \bar{a}]_{\mathbb{R}} [\sigma / \bar{b}] : \tau'' [\bar{\tau} / \bar{a}] [\sigma / \bar{b}] \sim_{\mathbb{R}} \tau'' [\bar{\tau}' / \bar{a}] [\sigma / \bar{b}]}$$

We can then conclude

$$\overline{\Gamma \vdash e' : \tau'' [\bar{\tau}' / \bar{a}] [\sigma / \bar{b}]}$$

as desired.

Putting this all together, we see that $\Gamma \vdash K \bar{\tau}' \bar{\sigma} \bar{\gamma}' \bar{e}' : D \bar{\tau}'$ as originally desired, and we are done. □

F Progress

We prove progress by first establishing that the global context is *consistent* (defined below). We do this by placing further restrictions on the global context and proving that these imply consistency. However, these restrictions are needed only for consistency, and it is possible to relax or change these in future versions of FC, as long as the consistency property holds by some mechanism.

F.1 Restrictions on axioms

There are two forms an axiom $C : [\bar{a}:\bar{\kappa}].\tau \sim_{\rho} \sigma$ can have, and different rules apply:

1. Newtype axioms: All of the following must hold.
 - (a) $\tau = N \bar{a}$
 - (b) $\rho = R$
 - (c) There must not be two axioms mentioning the same newtype N .
 - (d) The length of $\text{roles}(N)$ must match the arity of the axiom C .
2. Type family axioms: All of the following must hold.
 - (a) $\tau = F(\bar{\tau}')$
 - (b) $\rho = N$
 - (c) The types $\bar{\tau}'$ must not mention type families.
 - (d) Each $b \in \bar{a}$ must appear exactly once in the list $\bar{\tau}'$.
 - (e) Consider two axioms $C_1 : [\bar{a}:\bar{\kappa}].\tau_1 \sim_{\rho} \sigma_1$ and $C_2 : [\bar{b}:\bar{\kappa}'].\tau_2 \sim_{\rho} \sigma_2$ (where variables are renamed so that $\bar{a} \cap \bar{b} = \emptyset$). Then, if there exists some θ with $\theta(\tau_1) = \theta(\tau_2)$, it must be that $\theta(\sigma_1) = \theta(\sigma_2)$.

E.2 Consistency

Definition 35 (Type consistency). *Two types τ_1 and τ_2 are consistent if, whenever they are both value types:*

1. If $\tau_1 = H \bar{\sigma}$, then $\tau_2 = H \bar{\sigma}'$;
2. If $\tau_1 = \forall a:\kappa.\sigma$ then $\tau_2 = \forall a:\kappa.\sigma'$.

Note that if either τ_1 or τ_2 is *not* a value type (as defined in Appendix A.1), then they are vacuously consistent. Also, recall that a type headed by a **newtype** is not a value type.

Definition 36 (Context consistency). *The global context is consistent if, whenever $\bar{a}:\bar{\kappa} \vdash \gamma : \tau_1 \sim_R \tau_2$, τ_1 and τ_2 are consistent.*

In order to prove consistency, we define a nondeterministic type reduction relation $\tau \rightsquigarrow_\rho \sigma$, show that the relation preserves value type heads (when ρ is not phantom), and then show that any well-typed coercion corresponds to a path in the rewrite relation.

Here is the type rewrite relation:

$$\boxed{\tau \rightsquigarrow_\rho \sigma} \quad \text{Type reduction}$$

$$\frac{}{\tau \rightsquigarrow_\rho \tau} \text{ RED_REFL}$$

$$\frac{\tau_1 \rightsquigarrow_\rho \sigma_1 \quad \tau_2 \rightsquigarrow_{\mathbf{N}} \sigma_2}{\tau_1 \tau_2 \rightsquigarrow_\rho \sigma_1 \sigma_2} \text{ RED_APP}$$

$$\frac{\overline{\tau \rightsquigarrow_\rho \sigma} \quad \bar{\rho} \text{ is a prefix of } \text{roles}(H)}{H \bar{\tau} \rightsquigarrow_{\mathbf{R}} H \bar{\sigma}} \text{ RED_TYCONAPP}$$

$$\frac{\tau \rightsquigarrow_\rho \sigma}{\forall a : \kappa. \tau \rightsquigarrow_\rho \forall a : \kappa. \sigma} \text{ RED_FORALL}$$

$$\frac{\overline{\tau \rightsquigarrow_{\mathbf{N}} \sigma}}{F(\bar{\tau}) \rightsquigarrow_\rho F(\bar{\sigma})} \text{ RED_TYFAM}$$

$$\frac{C : [\bar{a} : \bar{\kappa}]. \tau_1 \sim_\rho \tau_2 \quad \rho \leq \rho'}{\tau_1[\sigma/a] \rightsquigarrow_{\rho'} \tau_2[\sigma/a]} \text{ RED_AXIOM}$$

$$\frac{}{\tau \rightsquigarrow_{\mathbf{P}} \sigma} \text{ RED_PHANTOM}$$

Lemma 37 (Simple rewrite substitution). *If $\tau_1 \rightsquigarrow_\rho \tau_2$, then $\tau_1[\sigma/a] \rightsquigarrow_\rho \tau_2[\sigma/a]$.*

Proof. By straightforward induction, noting that axioms have no free variables. \square

Lemma 38 (Rewrite substitution). *Let \bar{a} be the free variables in a type σ . If $\bar{a} : \bar{\rho} \vdash \sigma : \mathbf{R}$:*

1. *If $\overline{\tau \rightsquigarrow_\rho \tau'}$, then $\sigma[\tau/a] \rightsquigarrow_{\mathbf{R}} \sigma[\tau'/a]$;*
2. *If $\tau \rightsquigarrow_{\mathbf{N}} \tau'$, then $\sigma[\tau/a] \rightsquigarrow_{\mathbf{N}} \sigma[\tau'/a]$.*

Proof. Let $\Omega = \bar{a} : \bar{\rho}$. Proceed by induction on the structure of σ .

Case $\sigma = a$: There is thus only one free variable, a in σ . The one role ρ is \mathbf{R} . For clause (1), we know $\tau \rightsquigarrow_{\mathbf{R}} \tau'$, so we are done. For clause (2), we know $\tau \rightsquigarrow_{\mathbf{N}} \tau'$, so we are done.

Case $\sigma = \sigma_1 \sigma_2$:

Case (σ can be written as $H \bar{\sigma}$): Here, we assume that the length of $\bar{\sigma}$ is at most the length of $\text{roles}(H)$. If this is not the case, fall through to the “otherwise” case.

Clause (1): We know $\overline{\tau \rightsquigarrow_{\rho} \tau'}$. We must show that $H \overline{\sigma[\tau/a]} \rightsquigarrow_{\mathbf{R}} H \overline{\sigma[\tau'/a]}$. We will use RED_TYCONAPP. Let $\overline{\rho'}$ be a prefix of $\text{roles}(H)$ of the same length as $\overline{\sigma}$. We must show $\overline{\sigma[\tau/a]} \rightsquigarrow_{\rho'} \overline{\sigma[\tau'/a]}$.

Fix i . We will show that $\sigma_i[\tau/a] \rightsquigarrow_{\rho'_i} \sigma_i[\tau'/a]$.

Case ($\rho'_i = \mathbf{N}$): In order to use the induction hypothesis, we must show that for every j such that a_j appears free in σ_i , $\rho_j = \mathbf{N}$. To use Lemma 5, we must establish that $\Omega \vdash \sigma_i : \mathbf{N}$. We can get this by inversion on $\Omega \vdash H \overline{\sigma} : \mathbf{R}$ – whether by RTY_TYCONAPP or by RTY_APP, we get $\Omega \vdash \sigma_i : \mathbf{N}$. So, we can use the induction hypothesis and we are done.

Case ($\rho'_i = \mathbf{R}$): Inverting $\overline{a;\overline{\rho}} \vdash H \overline{\sigma} : \mathbf{R}$ gives us two possibilities:

Case RTY_TYCONAPP: Here, we see $\overline{\Omega \vdash \sigma} : \rho'$, and thus, that $\Omega \vdash \sigma_i : \mathbf{R}$ (because $\rho'_i = \mathbf{R}$). We can then use the induction hypothesis (and using Lemma 4 to make the contexts line up) and we are done.

Case RTY_APP: We invert repeatedly, and we either get $\Omega \vdash \sigma_i : \mathbf{N}$ or $\Omega \vdash \sigma_i : \rho'_i$, depending on whether we hit a RTY_TYCONAPP during the inversions. In the second case, we proceed as above (the RTY_TYCONAPP case). In the first case, we use Lemma 6 to conclude $\Omega \vdash \sigma_i : \mathbf{R}$ and use the induction hypothesis.

Case ($\rho'_i = \mathbf{P}$): We are done by RED_PHANTOM.

Clause (2): We know that $\overline{\tau \rightsquigarrow_{\mathbf{N}} \tau'}$. We must show that $H \overline{\sigma[\tau/a]} \rightsquigarrow_{\mathbf{N}} H \overline{\sigma[\tau'/a]}$. It is easier to consider the original type σ just as $\sigma_1 \sigma_2$, not as $H \overline{\sigma}$; fall through to the next case.

Otherwise:

Clause (1): We know $\overline{\tau \rightsquigarrow_{\rho} \tau'}$ and need to show that $(\sigma_1 \sigma_2)[\tau/a] \rightsquigarrow_{\mathbf{R}} (\sigma_1 \sigma_2)[\tau'/a]$.

The fact $\Omega \vdash \sigma_1 \sigma_2 : \mathbf{R}$ must be by RTY_APP. So, we can conclude $\Omega \vdash \sigma_1 : \mathbf{R}$ and $\Omega \vdash \sigma_2 : \mathbf{N}$. Then, we can use the induction hypothesis to get $\sigma_1[\tau/a] \rightsquigarrow_{\mathbf{R}} \sigma_1[\tau'/a]$.

To use the induction hypothesis for σ_2 , we must first establish that, for every j such that a_j appears free in σ_2 , $\tau_j \rightsquigarrow_{\mathbf{N}} \tau'_j$. Lemma 5 provides exactly this information, so we get $\overline{\sigma_2[\tau/a]} \rightsquigarrow_{\mathbf{N}} \overline{\sigma_2[\tau'/a]}$. We are done by RED_APP.

Clause (2): We know $\overline{\tau \rightsquigarrow_{\mathbf{N}} \tau'}$ and need to show that $(\sigma_1 \sigma_2)[\tau/a] \rightsquigarrow_{\mathbf{N}} (\sigma_1 \sigma_2)[\tau'/a]$.

We simply use induction to get:

$$\sigma_1[\tau/a] \rightsquigarrow_{\mathbf{N}} \sigma_1[\tau'/a]$$

$$\overline{\sigma_2[\tau/a]} \rightsquigarrow_{\mathbf{N}} \overline{\sigma_2[\tau'/a]}$$

We are done by RED_APP.

Case $\sigma = H$: We are done by RED_REFL.

Case $\sigma = \forall b:\kappa.\sigma'$: We assume that we have renamed variables so that $b \notin \overline{a}$. We see that inverting $\Omega \vdash \forall b:\kappa.\sigma' : \mathbf{R}$ gives us $\Omega, b:\mathbf{N} \vdash \sigma' : \mathbf{R}$, where \overline{a}, b are the free variables in σ' . We can then use the induction hypothesis and we are done by RED_FORALL.

Case $\sigma = F(\overline{\sigma})$: Inversion on $\Omega \vdash F(\overline{\sigma}) : \mathbf{R}$ gives us $\overline{\Omega \vdash \sigma} : \mathbf{N}$. We can then apply Lemma 5 to see that $\overline{\rho} = \mathbf{N}$. We then use the induction hypothesis repeatedly to

get

$$\overline{\sigma[\tau/a] \rightsquigarrow_{\mathbf{N}} \sigma[\tau'/a]}$$

We are now done by RED_TYFAM. □

Lemma 39 (Sub-roling in the rewrite relation). *If $\tau_1 \rightsquigarrow_{\mathbf{N}} \tau_2$, then $\tau_1 \rightsquigarrow_{\rho} \tau_2$.*

Proof. By straightforward induction on $\tau_1 \rightsquigarrow_{\mathbf{N}} \tau_2$. □

Lemma 40 (RED_APP/RED_TYCONAPP). *If $H \bar{\tau} \tau' \rightsquigarrow_{\mathbf{R}} H \bar{\sigma} \sigma'$ by RED_APP, the length of $\bar{\tau}$ is less than the length of $\text{roles}(H)$, then $H \bar{\tau} \tau' \rightsquigarrow_{\mathbf{R}} H \bar{\sigma} \sigma'$ also by RED_TYCONAPP.*

Proof. Fix H . We then proceed by induction on the length of $\bar{\tau}$.

Base case ($H \tau' \rightsquigarrow_{\mathbf{R}} H \sigma'$): The premises of RED_APP give us $H \rightsquigarrow_{\mathbf{R}} H$ and $\tau' \rightsquigarrow_{\mathbf{N}} \sigma'$.

Regardless of $\text{roles}(H)$, we can use the sub-roling lemma (Lemma 39) to show $\tau' \rightsquigarrow_{\rho} \sigma'$ and we are done. (In the case where $\text{roles}(H)$ is empty, an assumption is violated, and we are done anyway.)

Inductive case: Our inductive hypothesis says: if $H \bar{\tau} \rightsquigarrow_{\mathbf{R}} H \bar{\sigma}$ and $\tau' \rightsquigarrow_{\mathbf{N}} \sigma'$ (and the length of $\text{roles}(H)$ is sufficient), then $\bar{\tau} \rightsquigarrow_{\rho} \bar{\sigma}$ and $\tau' \rightsquigarrow_{\rho_i} \sigma'$, where $i = (\text{length of } \bar{\tau}) + 1$. We must show that, if $H \bar{\tau} \tau' \rightsquigarrow_{\mathbf{R}} H \bar{\sigma} \sigma'$ and $\tau'' \rightsquigarrow_{\mathbf{N}} \sigma''$ (and the length of $\text{roles}(H)$ is sufficient) then $\bar{\tau} \rightsquigarrow_{\rho} \bar{\sigma}$, $\tau' \rightsquigarrow_{\rho_i} \sigma'$, and $\tau'' \rightsquigarrow_{\rho_j} \sigma''$ (where $j = i + 1$).

Inverting $H \bar{\tau} \tau' \rightsquigarrow_{\mathbf{R}} H \bar{\sigma} \sigma'$ gives us several possibilities:

Case RED_REFL: We get $\bar{\tau} \rightsquigarrow_{\rho} \bar{\sigma}$ and $\tau' \rightsquigarrow_{\rho_i} \sigma'$ by RED_REFL. We get $\tau'' \rightsquigarrow_{\rho_j} \sigma''$ by Lemma 39.

Case RED_APP: We get our first two desiderata from use of the induction hypothesis and our last from Lemma 39.

Case RED_TYCONAPP: Our first two desiderata come from the premises of RED_TYCONAPP, and the last one comes from Lemma 39.

Case RED_AXIOM: This case is impossible, because there can be only one newtype axiom for a newtype, and its arity is greater than $(\text{length of } \bar{\tau}) + 1$. □

Lemma 41 (Pattern). *Let \bar{a} be the free variables in a type τ . We require that each variable a is mentioned exactly once in τ and that no type families appear in τ . Then, if, for some $\bar{\sigma}$, $\tau[\bar{\sigma}/\bar{a}] \rightsquigarrow_{\mathbf{N}} \tau'$, then there exist $\bar{\sigma}'$ such that $\tau' = \tau[\bar{\sigma}'/\bar{a}]$ and $\bar{\sigma} \rightsquigarrow_{\mathbf{N}} \bar{\sigma}'$.*

Proof. We proceed by induction on the structure of τ .

Case $\tau = a$: There is just one free variable (a), and thus just one type σ . We have $\sigma \rightsquigarrow_{\mathbf{N}} \tau'$. Let $\sigma' = \tau'$ and we are done.

Case $\tau = \tau_1 \tau_2$: Partition the free variables into a list \bar{b}_1 that appear in τ_1 and \bar{b}_2 that appear in τ_2 . This partition must be possible by assumption. Similarly, partition $\bar{\sigma}$ into $\bar{\sigma}_1$ and $\bar{\sigma}_2$. We can see that $\tau_1[\bar{\sigma}_1/\bar{b}_1] \tau_2[\bar{\sigma}_2/\bar{b}_2] \rightsquigarrow_{\mathbf{N}} \tau'$. Thus must be by RED_APP (noting that all newtype axioms are at role R). Thus, $\tau' = \tau'_1 \tau'_2$ and $\tau_1[\bar{\sigma}_1/\bar{b}_1] \rightsquigarrow_{\mathbf{N}} \tau'_1$ and $\tau_2[\bar{\sigma}_2/\bar{b}_2] \rightsquigarrow_{\mathbf{N}} \tau'_2$. We then use the induction hypothesis to get $\bar{\sigma}'_1$ and $\bar{\sigma}'_2$ such that $\tau'_1 = \tau_1[\bar{\sigma}'_1/\bar{b}_1]$ and $\tau'_2 = \tau_2[\bar{\sigma}'_2/\bar{b}_2]$. We conclude that $\bar{\sigma}'$ is the combination of $\bar{\sigma}'_1$ and $\bar{\sigma}'_2$, undoing the partition done earlier.

Case $\tau = H$: Trivial.

Case $\tau = \forall b:\kappa.\tau_0$: We first note that, according to the definition of \bar{a} , $b \notin \bar{a}$. We wish to use the induction hypothesis, but we must be careful because τ_0 may mention b multiple times. So, we linearise τ_0 into τ'_0 , replacing every occurrence of b with fresh variables \bar{b}' . (Note that \bar{b}' can be empty.) We know that $(\forall b:\kappa.\tau_0)[\bar{\sigma}/\bar{a}] \rightsquigarrow_{\mathbb{N}} \tau'$. We note that $(\forall b:\kappa.\tau_0)[\bar{\sigma}/\bar{a}] = \forall b:\kappa.(\tau_0[\bar{\sigma}/\bar{a}]) = \forall b:\kappa.(\tau'_0[\bar{\sigma}/\bar{a}][\bar{b}/\bar{b}'])$. (We have abused notation somewhat in the second substitution. There is only one b ; it is substituted for every variable in \bar{b}' .) Let $\bar{\sigma}''$ be $\bar{\sigma}$ appended with the right number of copies of b . Let \bar{a}' be \bar{a} appended with \bar{b}' . Then, we can say $\forall b:\kappa.(\tau'_0[\bar{\sigma}''/\bar{a}']) \rightsquigarrow_{\mathbb{N}} \tau'$. We invert to get that $\tau' = \forall b:\kappa.\tau''$ and $\tau'_0[\bar{\sigma}''/\bar{a}'] \rightsquigarrow_{\mathbb{N}} \tau''$. We can now use the induction hypothesis to get $\bar{\sigma}'''$ such that $\tau'' = \tau[\bar{\sigma}'''/\bar{a}']$ and $\bar{\sigma}'' \rightsquigarrow_{\mathbb{N}} \bar{\sigma}'''$. But, we can see that, b steps only to itself. Thus, the last entries in $\bar{\sigma}'''$ must be the same list of b s that $\bar{\sigma}''$ has. We let σ' be the prefix of $\bar{\sigma}'''$ without the b s, and we are done.

Case $\tau = F(\bar{\tau})$: Impossible, by assumption. □

Lemma 42 (Patterns). *Let \bar{a} be the free variables in a list of types $\bar{\tau}$. Assume each variable a is mentioned exactly once in $\bar{\tau}$ and that no type families appear in $\bar{\tau}$. If, for some $\bar{\sigma}$, $\tau[\bar{\sigma}/\bar{a}] \rightsquigarrow_{\mathbb{N}} \tau'$, then there exist $\bar{\sigma}'$ such that $\bar{\tau}' = \tau[\bar{\sigma}'/\bar{a}]$ and $\bar{\sigma} \rightsquigarrow_{\mathbb{N}} \bar{\sigma}'$.*

Proof. By induction on the length of $\bar{\tau}$.

Base case: Trivial.

Inductive case: We partition and recombine variables as in the $\tau_1 \tau_2$ case in the previous proof and proceed by induction. □

Lemma 43 (Local diamond). *If $\tau \rightsquigarrow_{\rho} \sigma_1$ and $\tau \rightsquigarrow_{\rho} \sigma_2$, then there exists σ_3 such that $\sigma_1 \rightsquigarrow_{\rho} \sigma_3$ and $\sigma_2 \rightsquigarrow_{\rho} \sigma_3$.*

Proof. If $\rho = P$, then the result is trivial, by RED_PHANTOM. So, we assume $\rho \neq P$.

If $\sigma_1 = \tau$ or $\sigma_2 = \tau$, the result is trivial. So, we assume that neither reduction is by RED_REFL.

By induction on the structure of τ :

Case $\tau = a$: We note that the left-hand side of an axiom can never be a bare variable, and so the only possibility of stepping is by RED_REFL. We are done.

Case $\tau = \tau_1 \tau_2$: Suppose $\rho = N$. All axioms at nominal role have a type family application on their left-hand side, so RED_AXIOM cannot apply. Thus, only RED_APP can be used, and we are done by induction.

Now, we can assume $\rho = R$. If $\tau_1 \tau_2$ cannot be rewritten as $H \bar{\tau}$ (for some H and some $\bar{\tau}$), then the only applicable rule is RED_APP (noting that relevant axiom left-hand sides can indeed be written as $H \bar{\tau}$) and we are done by induction.

So, we now rewrite τ as $H \bar{\tau}_0$. There are six possible choices of the two reductions, among RED_APP, RED_TYCONAPP, and RED_AXIOM. We handle each case separately:

Case RED_APP/RED_APP: We are done by induction.

Case RED_APP/RED_TYCONAPP: We apply Lemma 40 and finish by induction.

Case RED_APP/RED_AXIOM: Rewrite $\sigma_1 = \sigma_{11} \sigma_{12}$. We know then that $\tau_1 \rightsquigarrow_{\mathbb{R}} \sigma_{11}$ and $\tau_2 \rightsquigarrow_{\mathbb{N}} \sigma_{12}$. (Recall that $\tau_1 \tau_2 = \tau = H \bar{\tau}_0$.) We also know that $H \bar{\tau}_0 \rightsquigarrow_{\mathbb{R}} \sigma_2$ by a newtype axiom $C : [\bar{a}:\bar{\kappa}].H \bar{a} \sim_{\mathbb{R}} \sigma_0$, where $\sigma_2 = \sigma_0[\bar{\tau}_0/a]$.

By induction we can discover that σ_{11} has the form $H \bar{\sigma}$ – we know that τ_1 cannot reduce by RED_AXIOM because the restrictions on axioms say that newtype axioms are unique, and the axiom used on τ has a higher arity than any axiom that could be used on τ_1 . Thus, $\sigma_1 = H \bar{\sigma} \sigma_{12}$. The same axiom C applies here. Let $\bar{\sigma}' = \bar{\sigma}, \sigma_{12}$. So, we can step σ_1 to $\sigma_3 = \sigma_0[\bar{\sigma}'/a]$ by RED_AXIOM. Now, we must show $\sigma_2 \rightsquigarrow_{\mathbb{R}} \sigma_3$. We wish to apply the rewrite-substitution lemma (Lemma 38). We must show that $\bar{\tau}_0 \rightsquigarrow_{\rho} \bar{\sigma}'$, where $\bar{a}:\bar{\rho} \vdash \sigma_0 : \mathbb{R}$. This last fact is exactly what appears in the premise to ROLES_NEWTYPE (which, in turn, is guaranteed by the well-formedness of the global context). Now, we know $\tau = H \bar{\tau}_0$ and $\sigma_1 = H \bar{\sigma}'$, and that $\tau \rightsquigarrow_{\mathbb{R}} \sigma_1$ by RED_APP. We also know that an axiom is applicable to τ . Thus, the length of $\bar{\tau}$ must be the length of $\text{roles}(H)$, by context well-formedness. So, we can use Lemma 40 to get $\bar{\tau}_0 \rightsquigarrow_{\rho} \bar{\sigma}'$, as desired. We then apply Lemma 38 to conclude $\sigma_2 \rightsquigarrow_{\mathbb{R}} \sigma_3$, and we are done.

Case RED_TYCONAPP/RED_TYCONAPP: We are done by induction.

Case RED_TYCONAPP/RED_AXIOM: We see that $\sigma_1 = H \bar{\sigma}'$ where $\bar{\rho}$ is a prefix of $\text{roles}(H)$ and $\bar{\tau}_0 \rightsquigarrow_{\rho} \bar{\sigma}'$. We also see that $C : [\bar{a}:\bar{\kappa}].H \bar{a} \sim_{\mathbb{R}} \sigma_0$ and that $\sigma_2 = \sigma_0[\bar{\tau}_0/a]$.

Let $\sigma_3 = \sigma_0[\bar{\sigma}'/a]$. We can see that $\sigma_1 \rightsquigarrow_{\mathbb{R}} \sigma_3$ by RED_AXIOM. And, by Lemma 38 (the rewrite-substitution lemma), we see that $\sigma_2 \rightsquigarrow_{\mathbb{R}} \sigma_3$. So, we are done.

Case RED_AXIOM/RED_AXIOM: Consider the possibility that the two reductions are by different axioms. This would violate context well-formedness, so it is impossible. Thus, we can assume that the axiom used in both reductions is the same: $C : [\bar{a}:\bar{\kappa}].H \bar{a} \sim_{\mathbb{R}} \sigma_0$. The only way that σ_1 and σ_2 can be different is if the types substituted in the rule conclusion ($\bar{\sigma}$) are different in the two different reductions. Suppose then that we have $\bar{\sigma}$ and $\bar{\sigma}'$ so that $\sigma_1 = \sigma_0[\bar{\sigma}/a]$ and $\sigma_2 = \sigma_0[\bar{\sigma}'/a]$. It must be that $\tau = H \bar{\sigma}$ and that $\tau = H \bar{\sigma}'$. But, this tells us that $\bar{\sigma} = \bar{\sigma}'$ and thus that $\sigma_1 = \sigma_2$. We are done.

Case $\tau = H$: The only non-trivial step H can make is by RED_AXIOM. However, given that only one axiom for a newtype can exist, both steps must step to the same type, so we are done.

Case $\tau = \forall a:\kappa.\tau'$: We are done by induction.

Case $\tau = F(\bar{\tau})$: Here, two rules may apply. We handle the different possibilities separately:

Case RED_TYFAM/RED_TYFAM: We are done by induction.

Case RED_TYFAM/RED_AXIOM: Here, we know that $\sigma_1 = F(\bar{\sigma})$ where $\bar{\tau} \rightsquigarrow_{\mathbb{N}} \bar{\sigma}$, and that $\sigma_2 = \sigma_0[\bar{\sigma}'/a]$ where $C : [\bar{a}:\bar{\kappa}].F(\bar{\tau}') \sim_{\mathbb{N}} \sigma_0$ and $\bar{\tau} = \bar{\tau}'[\bar{\sigma}'/a]$.

We wish to use RED_AXIOM to reduce $F(\bar{\sigma})$. We apply Lemma 42 to get $\bar{\sigma}''$ such that $\bar{\sigma} = \bar{\tau}'[\bar{\sigma}''/a]$ and $\bar{\sigma}' \rightsquigarrow_{\mathbb{N}} \bar{\sigma}''$. We then use RED_AXIOM to get $\sigma_1 \rightsquigarrow_{\mathbb{N}} \sigma_3$,

where $\sigma_3 = \sigma_0[\overline{\sigma''/a}]$. Now, we must show that $\sigma_2 \rightsquigarrow_{\mathbf{N}} \sigma_3$. This comes directly from Lemma 38, and we are done.

Case RED_AXIOM/RED_AXIOM:

We have $C_1 : [\overline{a:\kappa}].F(\overline{\tau_1}) \sim_{\mathbf{N}} \sigma'_1$ and $C_2 : [\overline{b:\kappa'}].F(\overline{\tau_2}) \sim_{\mathbf{N}} \sigma'_2$. We also know that $\tau = F(\overline{\tau_1})[\overline{\sigma'/a}]$ and $\tau = F(\overline{\tau_2})[\overline{\sigma''/b}]$. Thus, $F(\overline{\tau_1})[\overline{\sigma'/a}] = F(\overline{\tau_2})[\overline{\sigma''/b}]$. Thus, $[\overline{\sigma'}, \overline{\sigma''/a}, \overline{b}]$ is a unifier for $F(\overline{\tau_1})$ and $F(\overline{\tau_2})$. Thus, by context well-formedness, we have $\sigma'_1[\overline{\sigma'/a}] = \sigma'_2[\overline{\sigma''/b}]$. But, $\sigma_1 = \sigma'_1[\overline{\sigma'/a}]$ and $\sigma_2 = \sigma'_2[\overline{\sigma''/b}]$, and so $\sigma_1 = \sigma_2$ and we are done. \square

Let the notation $\tau_1 \Leftrightarrow_{\rho} \tau_2$ mean that there exists a σ such that $\tau_1 \rightsquigarrow_{\rho}^* \sigma$ and $\tau_2 \rightsquigarrow_{\rho}^* \sigma$.

Lemma 44 (Confluence). *The rewrite relation \rightsquigarrow_{ρ} is confluent. That is, if $\tau \rightsquigarrow_{\rho}^* \sigma_1$ and $\tau \rightsquigarrow_{\rho}^* \sigma_2$, then $\sigma_1 \Leftrightarrow_{\rho} \sigma_2$.*

Proof. Confluence is a consequence of the local diamond property, Lemma 43. \square

Lemma 45 (Stepping preserves value type heads). *If τ_1 is a value type and $\tau_1 \rightsquigarrow_{\mathbf{R}} \tau_2$, then τ_2 has the same head as τ_1 .*

Proof. By induction, noting that the left-hand side of well-formed axioms are never value types. \square

Lemma 46 (Rewrite relation consistency). *If $\tau_1 \Leftrightarrow_{\mathbf{R}} \tau_2$, then τ_1 and τ_2 are consistent.*

Proof. If either τ_1 or τ_2 is not a value type, then we are trivially done. So, we assume τ_1 and τ_2 are value types. By assumption, there exists σ such that $\tau_1 \rightsquigarrow_{\mathbf{R}}^* \sigma$ and $\tau_2 \rightsquigarrow_{\mathbf{R}}^* \sigma$. By induction over the length of these reductions and the use of Lemma 45, we can see that σ must have the same head as both τ_1 and τ_2 . Thus, τ_1 and τ_2 have the same head, and are thus consistent. \square

Lemma 47 (Completeness of the rewrite relation). *If Γ binds no coercion variables and $\Gamma \vdash \gamma : \tau_1 \sim_{\rho} \tau_2$, then $\tau_1 \Leftrightarrow_{\rho} \tau_2$.*

Proof. By induction on $\Gamma \vdash \gamma : \tau_1 \sim_{\rho} \tau_2$.

Case CO_REFL: Trivial, as \Leftrightarrow_{ρ} is manifestly reflexive.

Case CO_SYM: By induction, as \Leftrightarrow_{ρ} is manifestly symmetric.

Case CO_TRANS: We adopt the variable names in the statement of the rule:

$$\frac{\Gamma \vdash \gamma_1 : \tau_1 \sim_{\rho} \tau_2 \quad \Gamma \vdash \gamma_2 : \tau_2 \sim_{\rho} \tau_3}{\Gamma \vdash \gamma_1 \circ \gamma_2 : \tau_1 \sim_{\rho} \tau_3} \text{ CO_TRANS}$$

By induction, we know $\tau_1 \Leftrightarrow_{\rho} \tau_2$ and $\tau_2 \Leftrightarrow_{\rho} \tau_3$. Thus, we must find σ_{13} such that $\tau_1 \rightsquigarrow_{\rho}^* \sigma_{13}$ and $\tau_3 \rightsquigarrow_{\rho}^* \sigma_{13}$. Note that there must be σ_{12} with $\tau_1 \rightsquigarrow_{\rho}^* \sigma_{12}$ and $\tau_2 \rightsquigarrow_{\rho}^* \sigma_{12}$, and there must be σ_{23} with $\tau_2 \rightsquigarrow_{\rho}^* \sigma_{23}$ and $\tau_3 \rightsquigarrow_{\rho}^* \sigma_{23}$. Thus, we can use Lemma 44 (confluence) to find a σ_{13} such that $\sigma_{12} \rightsquigarrow_{\rho}^* \sigma_{13}$ and $\sigma_{23} \rightsquigarrow_{\rho}^* \sigma_{13}$. By transitivity of $\rightsquigarrow_{\rho}^*$, we are done.

Case CO_TYCONAPP: We know by induction that $\bar{\tau} \Leftrightarrow_{\rho} \bar{\sigma}$. Let the list of common reducts be $\bar{\tau}'$. We can see that $H \bar{\tau} \rightsquigarrow_{\mathbf{R}}^* H \bar{\tau}'$ by repeated use of RED_TYCONAPP, and similarly for $H \bar{\sigma} \rightsquigarrow_{\mathbf{R}}^* H \bar{\tau}'$. Thus $H \bar{\tau}'$ is our common reduct and we are done.

Case CO_TYFAM: We are done by induction and repeated use of RED_TYFAM.

Case CO_APP: We are done by induction and repeated use of RED_APP.

Case CO_FORALL: We are done by induction and repeated use of RED_FORALL.

Case CO_PHANTOM: We are done by RED_PHANTOM.

Case CO_VAR: Not possible, as the context has no coercion variables.

Case CO_AXIOM: We are done by RED_AXIOM.

Case CO_NTH: We adopt the variable names in the rule:

$$\frac{\Gamma \vdash \gamma : H \bar{\tau} \sim_{\mathbf{R}} H \bar{\sigma} \quad \bar{\rho} \text{ is a prefix of } \text{roles}(H) \quad H \text{ is not a newtype}}{\Gamma \vdash \mathbf{nth}^i \gamma : \tau_i \sim_{\rho_i} \sigma_i} \text{ CO_NTH}$$

We know by induction that $H \bar{\tau} \Leftrightarrow_{\mathbf{R}} H \bar{\sigma}$. In other words, there exists some τ_0 such that $H \bar{\tau} \rightsquigarrow_{\mathbf{R}}^* \tau_0$ and $H \bar{\sigma} \rightsquigarrow_{\mathbf{R}}^* \tau_0$. We can see by induction on the number of steps in the derivation (and a nested induction in the RED_APP case) that τ_0 must have the form $H \bar{\tau}'$ for some $\bar{\tau}'$. In particular, note that no axioms can apply because H is not a newtype. Thus, each step is from either RED_APP or from RED_TYCONAPP. However, by Lemma 40, we can consider just the RED_TYCONAPP case. This says that $\tau_i \rightsquigarrow_{\rho_i}^* \tau'_i$ and $\sigma_i \rightsquigarrow_{\rho_i}^* \tau'_i$, as desired, so we are done.

Case CO_LEFT: We adopt the variable names from the rule:

$$\frac{\Gamma \vdash \gamma : \tau_1 \tau_2 \sim_{\mathbf{N}} \sigma_1 \sigma_2 \quad \Gamma \vdash \tau_1 : \kappa \quad \Gamma \vdash \sigma_1 : \kappa}{\Gamma \vdash \mathbf{left} \gamma : \tau_1 \sim_{\mathbf{N}} \sigma_1} \text{ CO_LEFT}$$

We know by induction that $\tau_1 \tau_2 \Leftrightarrow_{\mathbf{N}} \sigma_1 \sigma_2$. The steps to reach the common reduct must all be RED_APP, because newtype axioms are all at role R. Thus, the common reduct must be $\tau'_1 \tau'_2$ where $\tau_1 \rightsquigarrow_{\mathbf{N}}^* \tau'_1$, and $\sigma_1 \rightsquigarrow_{\mathbf{N}}^* \tau'_1$, so we are done.

Case CO_RIGHT: Similar to previous case.

Case CO_INST: We adopt the variable names from the rule:

$$\frac{\Gamma \vdash \gamma : \forall a : \kappa. \tau_1 \sim_{\rho} \forall a : \kappa. \sigma_1 \quad \Gamma \vdash \tau : \kappa}{\Gamma \vdash \gamma @ \tau : \tau_1[\tau/a] \sim_{\rho} \sigma_1[\tau/a]} \text{ CO_INST}$$

We know by induction that $\forall a : \kappa. \tau_1 \Leftrightarrow_{\rho} \forall a : \kappa. \sigma_1$. We can easily see by inspection of the rewrite relation that the common reduct must have the form $\forall a : \kappa. \tau_0$ for some τ_0 . We can also see by a straightforward induction that $\tau_1 \rightsquigarrow_{\rho}^* \tau_0$ and $\sigma_1 \rightsquigarrow_{\rho}^* \tau_0$. We must show that $\tau_1[\tau/a] \rightsquigarrow_{\rho}^* \tau_0[\tau/a]$ and $\sigma_1[\tau/a] \rightsquigarrow_{\rho}^* \tau_0[\tau/a]$. These facts come from an induction over the lengths of the derivations and the use of the simple rewrite substitution lemma, Lemma 37.

Case CO_SUB: We adopt the variable names in the rule:

$$\frac{\Gamma \vdash \gamma : \tau \sim_{\mathbf{N}} \sigma}{\Gamma \vdash \mathbf{sub} \gamma : \tau \sim_{\mathbf{R}} \sigma} \text{ CO_SUB}$$

We know that $\tau \Leftrightarrow_{\mathbb{N}} \sigma$ and we need $\tau \Leftrightarrow_{\mathbb{R}} \sigma$. This follows by induction over the lengths of the reduction and the use of Lemma 39. \square

Lemma 48 (Consistency). *The global context is consistent.*

Proof. Take a γ such that $\bar{a}:\bar{\kappa} \vdash \gamma : \tau_1 \sim_{\mathbb{R}} \tau_2$. By the completeness of the rewrite relation (Lemma 47), we see that $\tau_1 \Leftrightarrow_{\mathbb{R}} \tau_2$. But, the rewrite relation consistency lemma (Lemma 46) tells us that τ_1 and τ_2 are consistent. Thus, the context admits only consistent coercions and is itself consistent. \square

F.3 Progress

Lemma 49 (Canonical forms).

1. If $\Gamma \vdash v : \tau_1 \rightarrow \tau_2$, then v is either $\lambda x:\tau_1.e'$ or $K\bar{\tau}\bar{\gamma}\bar{e}$.
2. If $\Gamma \vdash v : \forall a:\kappa.\tau$, then v is either $\Lambda a:\kappa.v'$ or $K\bar{\tau}$.
3. If $\Gamma \vdash v : \phi \Rightarrow \tau$, then v is either $\lambda c:\phi.e'$ or $K\bar{\tau}\bar{\gamma}$.
4. If $\Gamma \vdash v : D\bar{\sigma}$, then v is $K\bar{\tau}\bar{\gamma}\bar{e}$.

Note that, following Haskell's lazy semantics, data constructors do not evaluate their arguments before pattern matching.

Lemma 50 (Value types). *If $\Gamma \vdash v : \tau$, then τ is a value type.*

Proof. If v is an abstraction, then the result is trivial. So, we assume that $v = K\bar{\tau}\bar{\gamma}\bar{e}$. Induction on the lengths of the lists of arguments yields

$$K : \forall \bar{a}:\bar{\kappa}.\forall \bar{b}:\bar{\kappa}'.\bar{\phi} \Rightarrow \bar{\sigma} \rightarrow D\bar{a}$$

We can see (again, by induction on the argument lists) that no matter what K is applied to, its type will always be a value type, headed by one of \forall , \Rightarrow , \rightarrow or D , all of which form value types. \square

Theorem 51 (Progress). *If $\bar{a}:\bar{\kappa} \vdash e : \tau$, then either e is a value or a coerced value, or $e \xrightarrow{\bar{a}:\bar{\kappa}} e'$ for some e' .*

Note that, unlike most proofs of progress, here we allow type variables in the context. This is necessary to deal with evaluation under Λ .

Proof. We proceed by induction on the typing judgement $\bar{a}:\bar{\kappa} \vdash e : \tau$.

Case TM_VAR: Cannot happen in a context without term variables.

Case TM_ABS: Trivial.

Case TM_APP: We know $e = e_1 e_2$. By induction, we know that e_1 is either a value, a coerced value, or steps to e'_1 . If e_1 steps, we are done by S_APP_CONG. If e_1 is a value, the canonical forms lemma now gives us several cases:

Case $e_1 = \lambda x:\tau.e_3$: We are done by S_BETA.

Case $e_1 = K\bar{\tau}\bar{\gamma}\bar{e}$: Then, $e_1 e_2$ is a value.

If e_1 is a coerced value $v \triangleright \gamma$, then by the value types lemma (Lemma 50) and the consistency lemma (Lemma 48), the type of v must be headed by (\rightarrow) . We are done by S_PUSH .

Case TM_TABS : We know $e = \Lambda a:\kappa.e_1$. By induction, either e_1 steps to e'_1 , or e_1 is a value or a coerced value. If e_1 steps to e'_1 , we are done by S_TABS_CONG .

Otherwise if e_1 is a value, then $\Lambda a:\kappa.e_1$ is also a value. If e_1 is a coerced value $v \triangleright \gamma$, then we are done by S_APUSH .

Case TM_TAPP : Similar to TM_APP case.

Case TM_CABS : Trivial.

Case TM_CAPP : Similar to TM_APP case.

Case $TM_DATACON$: e is a value.

Case TM_CASE : We adopt the variable names from the rule:

$$\frac{\begin{array}{l} \Gamma \vdash e : D \bar{\sigma} \\ \Gamma \vdash \tau : \star \\ \forall alt_i \text{ s.t. } alt_i \in \overline{alt}, \\ alt_i = K_i \bar{a}_i \bar{c}_i \bar{x}_i \rightarrow e_i \\ K_i : \forall \bar{a}'_i : \kappa_i. \forall \bar{b}'_i : \kappa'_i. \bar{\phi}_i \Rightarrow \bar{\tau}_i \rightarrow D \bar{a}'_i \\ \Gamma, \bar{a}_i : \kappa'_i, (\bar{c}_i : \bar{\phi}_i, \bar{x}_i : \bar{\tau}_i) [\bar{\sigma} / \bar{a}'_i] [a_i / \bar{b}'_i] \vdash e_i : \tau \\ \overline{alt} \text{ is exhaustive} \end{array}}{\Gamma \vdash \mathbf{case}_\tau e \text{ of } \overline{alt} : \tau} \quad TM_CASE$$

We know by induction that e is a value, a coerced value, or $e \xrightarrow{\bar{a}:\bar{\kappa}} e'$ for some e' .

If e steps, then we are done by S_CASE_CONG .

We see that e has a value type. Therefore (Lemmas 50 and 48), if it has the form $v \triangleright \gamma$, the value v has a type headed by Δ as well. Thus (Lemma 49) $v = K \bar{\tau} \bar{\gamma} \bar{e}$ and we apply S_KPUSH , noting that the premises are all satisfied by straightforward use of typing judgements.

The final case is that e is a value. By the canonical forms lemma, we see that $e = K \bar{\tau} \bar{\gamma} \bar{e}$. Thus, S_IOTA applies, noting that the match must be exhaustive.

Case TM_CAST : We adopt the variable names from the rule:

$$\frac{\begin{array}{l} \Gamma \vdash e : \tau_1 \\ \Gamma \vdash \gamma : \tau_1 \sim_R \tau_2 \end{array}}{\Gamma \vdash e \triangleright \gamma : \tau_2} \quad TM_CAST$$

By induction, we know that e is a value, a coerced value, or $e \xrightarrow{\bar{a}:\bar{\kappa}} e'$.

If e steps, we are done by S_CAST_CONG .

If e is a value, then $e \triangleright \gamma$ is a coerced value, and we are done.

If e is a coerced value, then we are done by S_TRANS .

Case TM_CONTRA : We adopt the variable names from the rule:

$$\frac{\begin{array}{l} \emptyset \vdash \gamma : H_1 \sim_N H_2 \quad H_1 \neq H_2 \\ \Gamma \vdash \tau : \star \end{array}}{\Gamma \vdash \mathbf{contra} \gamma \tau : \tau} \quad TM_CONTRA$$

By completeness of the rewrite relation (Lemma 47), we know that $H_1 \Leftrightarrow_N H_2$. But, if $H \rightsquigarrow_N H'$, then $H = H'$ (by induction on $H \rightsquigarrow_N H'$, noting that all newtype

axioms are at role R). So $H_1 = H_2$, contradicting a premise to this rule. Thus, this case cannot happen. \square

G Role inference

Lemma 52 (Walking). *Let \bar{a} be the parameters to some type constant T . For some type σ , let \bar{b} be the free variables in σ that are not in \bar{a} . Let $\bar{\rho}$ be a list of roles of the same length as \bar{a} . Let $\Omega = \bar{a}; \bar{\rho}; \bar{b}; \mathbb{N}$.*

If $\text{walk}(T, \sigma)$ makes no change to the role of any of the \bar{a} , then $\Omega \vdash \sigma : \mathbb{R}$.

Proof. By induction on the structure of σ :

Case $\sigma = a'$: By assumption, it must be that $a' : \mathbb{R} \in \Omega$ or $a' : \mathbb{N} \in \Omega$. In either case, we can derive $\Omega \vdash a' : \mathbb{R}$, so we are done.

Case $\sigma = \sigma_1 \sigma_2$: We check if σ can also be written as $H' \bar{\tau}$.

Case $\sigma = H' \bar{\tau}$: Let $\bar{\rho}' = \text{roles}(H')$. In order to conclude $\Omega \vdash H' \bar{\tau} : \mathbb{R}$, we will show that $\bar{\Omega} \vdash \bar{\tau} : \bar{\rho}'$. Fix i ; we will show $\Omega \vdash \tau_i : \rho'_i$. Here, we have three cases:

Case $\rho'_i = \mathbb{N}$: By assumption, it must be that all the free variables in τ_i are assigned to \mathbb{N} in Ω . Thus, by Lemma 5, we have $\Omega \vdash \tau_i : \mathbb{N}$ and we are done.

Case $\rho'_i = \mathbb{R}$: By assumption, it must be that $\text{walk}(T, \tau_i)$ makes no change. We then use the induction hypothesis to say that $\Omega \vdash \tau_i : \mathbb{R}$, and we are done.

Case $\rho'_i = \mathbb{P}$: We are done by `RTY_PHANTOM`.

Other applications: We wish to use `RTY_APP`. Thus, we must show that $\Omega \vdash \sigma_1 : \mathbb{R}$ and $\Omega \vdash \sigma_2 : \mathbb{N}$. For the former, we see that $\text{walk}(T, \sigma_1)$ must make no change, and we are done by induction. For the latter, we see that all the free variables in σ_2 must be assigned to \mathbb{N} , and we are done by Lemma 5.

Case $\sigma = H$: We are done by immediate application of `RTY_TYCONAPP`.

Case $\sigma = \forall a' : \kappa. \sigma_1$: We are done by induction, noting that in `RTY_FORALL`, a' gets assigned role \mathbb{N} when checking σ_1 . This matches our expectations that the type variables \bar{b} are at role \mathbb{N} in the inductive hypothesis.

Case $\sigma = F(\bar{\tau})$: Repeated use of Lemma 5 tells us that $\bar{\Omega} \vdash \bar{\tau} : \bar{\mathbb{N}}$. We are done by `RTY_TYFAM`. \square

Theorem 53. *The role inference algorithm always terminates.*

Proof. First, we observe that the walk procedure always terminates, as it is structurally recursive.

For the algorithm to loop in step 4, a role assigned to a variable must have changed. Yet, there are a finite number of such variables, and each variable may be updated only at most twice (from \mathbb{P} to \mathbb{R} and from \mathbb{R} to \mathbb{N}). Thus, at some point no more updates will happen and the algorithm will terminate. \square

Theorem 54 (Role inference is sound). *After running the role inference algorithm, $\text{roles}(H) \models H$ will hold for all H .*

Proof. We handle the data type case first. Fix a D . We will show that $\text{roles}(D) \models D$. Because the role inference algorithm has terminated, we know that $\text{walk}(D, \sigma)$ has caused no change for every σ that appears as a coercion type or term-level argument type in a constructor for D . Choose a constructor K , such that

$$K : \forall \bar{a}:\bar{\kappa}.\forall \bar{b}:\bar{\kappa}'.\bar{\phi} \Rightarrow \bar{\sigma} \rightarrow D \bar{a}$$

Let $\bar{\rho} = \text{roles}(D)$ and $\Omega = \bar{a}:\bar{\rho}, \bar{b}:\bar{N}$. We have satisfied the premises of the walking lemma (Lemma 52), and thus we can conclude that $\Omega \vdash \sigma : R$. We have shown $\text{roles}(D) \models D$ by ROLES_DATA.

The newtype case is similar, using the right-hand side of the newtype definition in place of σ . \square

Lemma 55 (Stumbling). *Let \bar{a} be the parameters to some type constant T . For some type σ , let \bar{b} be the free variables in σ that are not in \bar{a} . Let $\bar{\rho}$ be a list of roles of the same length as \bar{a} . Let $\Omega = \bar{a}:\bar{\rho}, \bar{b}:\bar{N}$.*

If $\text{walk}(T, \sigma)$ were modified to skip one of its attempts to mark a variable, then it is not possible to conclude $\Omega \vdash \sigma : R$.

Proof. By induction on the structure of σ :

Case $\sigma = a'$: If that mark were not done, then Ω would contain $a':P$; this clearly violates $\Omega \vdash a' : R$.

Case $\sigma = \sigma_1 \sigma_2$: We check if σ can also be written as $H' \bar{\tau}$.

Case $\sigma = H' \bar{\tau}$: Let $\bar{\rho}' = \text{roles}(H')$. Fix i .

Case $\rho'_i = N$: If we do not mark every free variable in τ_i as N , then it would be impossible to conclude $\Omega \vdash \tau_i : N$, by Lemma 5. Thus, we would not be able to conclude $\Omega \vdash H' \bar{\tau} : R$ by RTY_TYCONAPP. What about by RTY_APP? This, too, would require $\Omega \vdash \tau_i : N$, which we are unable to do.

Case $\rho'_i = R$: By induction, it is not possible to conclude $\Omega \vdash \tau_i : R$, and thus impossible to use RTY_TYCONAPP. What about RTY_APP? This would require $\Omega \vdash \tau_i : N$, which is not possible via the contrapositive of Lemma 6.

Case $\rho'_i = P$: There is no marking to be done here, so the assumption that walk is modified is false.

Other applications: Suppose the skipped marking were in the recursive call. Then, by induction, it is not possible to conclude $\Omega \vdash \sigma_1 : R$. Thus, it is not possible to conclude $\Omega \vdash \sigma_1 \sigma_2 : R$ by RTY_APP.

Now, suppose the skipped marking is when marking all free variables in σ_2 as N . In this case, we know that $\Omega \vdash \sigma_2 : N$ is impossible (by Lemma 5) and thus we cannot use RTY_APP.

Case $\sigma = H$: No mark was skipped, so the assumption that walk is modified is false.

Case $\sigma = \forall a':\kappa.\sigma_1$: We are done by induction, noting that in RTY_FORALL, a' gets assigned role N when checking σ_1 . This matches our expectations that the type variables \bar{b} are at role N in the inductive hypothesis.

Case $\sigma = F(\bar{\tau})$: If one of the variables free in the $\bar{\tau}$ were not marked as \mathbb{N} , then it would be impossible to conclude $\Omega \vdash \tau_i : \mathbb{N}$ for that τ_i (by Lemma 5). Thus, we would be unable to use `RTY_TYFAM`.

□

Theorem 56 (Role inference is optimal). *Suppose H has no role annotation. After running the role inference algorithm, any loosening of the roles assigned to H (a change from ρ to ρ' , where $\rho \leq \rho'$ and $\rho \neq \rho'$) would violate $\text{roles}(H) \models H$.*

Proof. Every time the role inference algorithm changes an assigned role from ρ' to ρ , it is the case that $\rho \leq \rho'$ and $\rho \neq \rho'$. Thus, all we must show is that every change the algorithm makes is necessary – that is, not making the change would then violate $\text{roles}(H) \models H$.

Role inference runs only on algebraic data types, so we need only concern ourselves with T s, not general H s. In both the data type and newtype cases, showing $\text{roles}(T) \models T$ requires showing $\Omega \vdash \sigma : R$, where $\Omega = \bar{a}:\bar{\rho}, \bar{b}:\mathbb{N}$ and \bar{a} are the parameters to T and \bar{b} are the remaining free variables of σ . (In the newtype case, \bar{b} is empty.) The list of roles $\bar{\rho}$ is $\text{roles}(T)$. So, we must show that skipping any change in the `walk(T, σ)` algorithm means that $\Omega \vdash \sigma : R$ would not be derivable. This is precisely what Lemma 55 shows and so we are done. □

Theorem 57 (Role annotations only tighten roles). *Suppose a role annotation assigns roles $\bar{\rho}$ to H . If roles $\bar{\rho}'$ were inferred for a definition H' identical to H but missing H 's role annotation, then $\bar{\rho} \leq \bar{\rho}'$.*

Proof. By Theorem 54, we know $\text{roles}(H) \models H$. Yet by Theorem 56, we know that any loosening of the roles of H' would violate $\text{roles}(H') \models H'$. The \models judgement does not consult role annotations; thus $\bar{\rho} \models H$ implies $\bar{\rho} \models H'$. If $\text{roles}(H)$ were looser than $\text{roles}(H')$, we can still derive $\text{roles}(H') \models H'$ thus leading to a contradiction. □

We wish to prove a *principal roles* property, stating that a unique “best” (most permissive) role assignment exists. To do this, we must consider multiple different values of $\text{roles}(H)$ for a given H . We thus introduce *role assignment environments* Ψ :

$$\Psi ::= \emptyset \mid \Psi, H : \bar{\rho}$$

In effect, the *roles* operator we use elsewhere is an implicit, global role assignment environment.

We also introduce the notation *sound* Ψ to mean $\forall H \in \Psi, \Psi(H) \models H$.

Finally, we parameterised the judgement $\Omega \vdash \tau : \rho$ by a role assignment environment Ψ , writing $\Psi; \Omega \vdash \tau : \rho$. The only rule that changes is `RTY_TYCONAPP`, which now looks like this:

$$\frac{\bar{\rho} \text{ is a prefix of } \Psi(H) \quad \overline{\Psi; \Omega \vdash \tau : \rho}}{\Psi; \Omega \vdash H \bar{\tau} : R} \text{ RTY_TYCONAPP}$$

Lemma 58 (Roles of an applied type constant). *If $\Psi; \Omega \vdash H \bar{\tau} : R$ and $\bar{\rho}$ is a prefix of $\Psi(H)$, then $\overline{\Psi; \Omega \vdash \tau : \rho}$.*

Proof. Proceed by induction on the length of $\bar{\tau}$.

Case $\bar{\tau} = \emptyset$: The conclusion is trivial. We are done.

Case $\bar{\tau} = \bar{\sigma}, \sigma_0$: Inversion on $\Psi; \Omega \vdash H \bar{\sigma} \sigma_0 : R$ gives us two cases:

Case RTY_TYCONAPP: The result is immediate from the premises of RTY_TYCONAPP.

Case RTY_APP: We know $\Psi; \Omega \vdash H \bar{\sigma} : R$ and $\Psi; \Omega \vdash \sigma_0 : N$. Lemma 6 tells us that $\Psi; \Omega \vdash \sigma_0 : \rho$ for any ρ . We are thus done by a use of the induction hypothesis. \square

Lemma 59 (Maximising roles). *If $\Psi_1; \Omega_1 \vdash \tau : \rho_1$ and $\Psi_2; \Omega_2 \vdash \tau : \rho_2$, then $\max(\Psi_1, \Psi_2); \max(\Omega_1, \Omega_2) \vdash \tau : \max(\rho_1, \rho_2)$.*

Proof. If either ρ_1 or ρ_2 is P, then $\max(\rho_1, \rho_2) = P$ and we are done by RTY_PHANTOM. We thus assume that neither ρ_1 nor ρ_2 is P. Proceed by induction on the structure of τ :

Case $\tau = a$: Inverting RTY_VAR, we must have $a; \rho'_1 \in \Omega_1$ and $a; \rho'_2 \in \Omega_2$ for some $\rho'_1 \leq \rho_1$ and $\rho'_2 \leq \rho_2$. It must be that $\max(\rho'_1, \rho'_2) \leq \max(\rho_1, \rho_2)$, and so we are done with this case.

Case $\tau = \tau_1 \tau_2$: We now have several cases, depending on the inversions of $\Psi_1; \Omega_1 \vdash \tau_1 \tau_2 : \rho_1$ and $\Psi_2; \Omega_2 \vdash \tau_1 \tau_2 : \rho_2$:

Case RTY_TYCONAPP/RTY_TYCONAPP: We adopt the metavariable names from the rule, as written above. If $\bar{\rho}_1$ is a prefix of $\Psi_1(H)$ and $\bar{\rho}_2$ is a prefix of $\Psi_2(H)$, then we know $\Psi_i; \Omega_i \vdash \tau : \rho_i$. Choose a specific $\tau \in \bar{\tau}$. We must prove $\max(\Psi_1, \Psi_2); \max(\Omega_1, \Omega_2) \vdash \tau : \max(\rho_1, \rho_2)$. We are done by the induction hypothesis.

Case RTY_TYCONAPP/RTY_APP: In this case, we know that $\tau = H \bar{\sigma} \sigma_0$. We need to prove $\max(\Psi_1, \Psi_2); \max(\Omega_1, \Omega_2) \vdash \tau : R$. (We know the role after the colon is R because of inversion on RTY_TYCONAPP.) If $\bar{\rho}_1, \rho_{10}$ is a prefix of $\Psi_1(H)$, then we know $\Psi_1; \Omega_1 \vdash \sigma : \rho_1$, $\Psi_1; \Omega_1 \vdash \sigma_0 : \rho_{10}$, $\Psi_2; \Omega_2 \vdash H \bar{\sigma} : R$, and $\Psi_2; \Omega_2 \vdash \sigma_0 : N$. Lemma 58 tells us that $\Psi_2; \Omega_2 \vdash \sigma : \rho_2$ (with $\bar{\rho}_2, \rho_{20}$ a prefix of $\Psi_2(H)$). Let $\bar{\rho}_3 = \max(\bar{\rho}_1, \bar{\rho}_2)$, where the maximum is computed pointwise. The induction hypothesis tells us $\max(\Psi_1, \Psi_2); \max(\Omega_1, \Omega_2) \vdash \sigma : \rho_3$. Lemma 6 tells us that $\Psi_2; \Omega_2 \vdash \sigma_0 : \rho_{20}$ and thus (by the induction hypothesis) that

$$\max(\Psi_1, \Psi_2); \max(\Omega_1, \Omega_2) \vdash \sigma_0 : \max(\rho_{10}, \rho_{20}).$$

We can see that $\bar{\rho}_3, \rho_{30}$ is a prefix of $\max(\Psi_1, \Psi_2)(H)$ and we are thus done by RTY_TYCONAPP.

Case RTY_APP/RTY_APP: Here, we know the following:

- $\Psi_1; \Omega_1 \vdash \tau_1 : \rho_1$
- $\Psi_1; \Omega_1 \vdash \tau_2 : N$
- $\Psi_2; \Omega_2 \vdash \tau_1 : \rho_2$
- $\Psi_2; \Omega_2 \vdash \tau_2 : N$

The induction hypothesis gives us $\max(\Psi_1, \Psi_2); \max(\Omega_1, \Omega_2) \vdash \tau_1 : \max(\rho_1, \rho_2)$ and $\max(\Psi_1, \Psi_2); \max(\Omega_1, \Omega_2) \vdash \tau_2 : N$. We are done by RTY_APP.

Case $\tau = H$: We are done by Lemma 8.

Case $\tau = \forall a:\kappa.\tau_0$: Inversion on RTY_FORALL tells us $\Psi_1;\Omega_1,a:\mathbb{N} \vdash \tau_0:\rho_1$ and $\Psi_2;\Omega_2,a:\mathbb{N} \vdash \tau_0:\rho_2$. The induction hypothesis gives us $\max(\Psi_1, \Psi_2); \max(\Omega_1, a:\mathbb{N}, \Omega_2, a:\mathbb{N}) \vdash \tau_0:\max(\rho_1, \rho_2)$. It is easy to see this is equivalent to $\max(\Psi_1, \Psi_2); \max(\Omega_1, \Omega_2), a:\mathbb{N} \vdash \tau_0:\max(\rho_1, \rho_2)$ and so we are done by RTY_FORALL.

Case $\tau = F(\bar{\sigma})$: By induction. □

Lemma 60 (Maximal role assignment environments). *Assume a given list of constants \bar{H} with role assignment environments Ψ_1 and Ψ_2 , both defined over \bar{H} . If sound Ψ_1 and sound Ψ_2 , then sound $\max(\Psi_1, \Psi_2)$.*

Proof. Let $\Psi_0 = \max(\Psi_1, \Psi_2)$.

Without loss of generality, choose a specific type constant $H \in \bar{H}$. We must show $\Psi_0(H) \models H$ given $\Psi_1(H) \models H$ and $\Psi_2(H) \models H$. Let $\bar{\rho}_i = \Psi_i(H)$ for $i = 0, 1, 2$. By definition of Ψ_0 , $\bar{\rho}_0 = \max(\bar{\rho}_1, \bar{\rho}_2)$.

Proceed by case analysis on H :

Case $H = D$: We must proceed by ROLES_DATA:

$$\frac{\begin{array}{l} \forall \bar{a}, \bar{b}, \bar{\sigma} \text{ s.t. } K : \forall \bar{a}:\bar{\kappa}. \forall \bar{b}:\bar{\kappa}'. \bar{\phi} \Rightarrow \bar{\sigma} \rightarrow D \bar{a}, \\ \forall \tau \text{ s.t. } \tau \in \bar{\sigma} \vee \tau \in \bar{\phi}, \\ \bar{a}:\bar{\rho}, \bar{b}:\mathbb{N} \vdash \tau : R \end{array}}{\bar{\rho} \models D} \text{ ROLES_DATA}$$

Fix a particular τ as chosen by the premise of ROLES_DATA. We know both that $\Psi_1;\bar{a}:\bar{\rho}_1, \bar{b}:\mathbb{N} \vdash \tau : R$ and that $\Psi_2;\bar{a}:\bar{\rho}_2, \bar{b}:\mathbb{N} \vdash \tau : R$. We can conclude that $\Psi_0;\bar{a}:\bar{\rho}_0, \bar{b}:\mathbb{N} \vdash \tau : R$ by Lemma 59 and are done with this case.

Case $H = N$: Inverting ROLES_NEWTYPE,

$$\frac{C : [\bar{a}:\bar{\kappa}]. N \bar{a} \sim_R \sigma \quad \bar{a}:\bar{\rho} \vdash \sigma : R}{\bar{\rho} \models N} \text{ ROLES_NEWTYPE}$$

we get $\Psi_1;\bar{a}:\bar{\rho}_1 \vdash \sigma : R$ and $\Psi_2;\bar{a}:\bar{\rho}_2 \vdash \sigma : R$. Use Lemma 59 to get $\Psi_0;\bar{a}:\bar{\rho}_0 \vdash \sigma : R$ and we are done by ROLES_NEWTYPE.

Case $H = (\rightarrow)$: Since $\Psi_1((\rightarrow)) \models (\rightarrow)$ and $\Psi_2((\rightarrow)) \models (\rightarrow)$, it must be that $\Psi_1((\rightarrow)) = \Psi_2((\rightarrow)) = R, R$. Thus $\Psi_0((\rightarrow)) = R, R$ and we are done by ROLES_ARROW.

Case $H = (\Rightarrow)$: Similar to previous case.

Case $H = (\sim_\rho)$: Similar to previous case. □

Theorem 61 (Principal role assignments). *For a given set of type constants \bar{H} , there is at most one choice of role assignments $\text{roles}(H)$ that is optimal and such that $\text{roles}(H) \models H$.*

Proof. With the concept of role assignment environments Ψ at our disposal, we can restate this theorem: Given Ψ_1 and Ψ_2 such that both are optimal and both are sound, it must be that $\Psi_1 = \Psi_2$.

We prove by contradiction. Suppose we have optimal, sound Ψ_1 and Ψ_2 such that $\Psi_1 \neq \Psi_2$. Lemma 60 tells us that $\max(\Psi_1, \Psi_2)$ is sound. But if $\max(\Psi_1, \Psi_2)$ differs from Ψ_i (for $i = 1, 2$), then Ψ_i is not optimal, violating our assumption. □

H Type erasure

In order to show that coercions are zero-cost, we prove a type erasure property, saying that evaluation of an FC expression simulates the evaluation of an expression in a simpler, erased language, which is devoid of types and coercions. It is this erased language that is actually evaluated at runtime. As it contains no coercions, we show that coercions truly are zero-cost.

The definition of the erased language is as follows:

o	$::=$	$x \mid \lambda x.o \mid \lambda \bullet.o \mid o_1 o_2 \mid o \bullet \mid K \mid \mathbf{case} \, o \, \mathbf{of} \, \overline{ealt}$	erased expressions
$ealt$	$::=$	$K\bar{x} \rightarrow o$	erased case alternative
w	$::=$	$\lambda x.o \mid K\bar{o}$	erased values

An erased expression steps according to the following small-step semantics:

$$\boxed{o \Longrightarrow o'}$$

$$\frac{}{(\lambda x.o_1) o_2 \Longrightarrow o_1[o_2/x]} \text{ E_BETA}$$

$$\frac{}{(\lambda \bullet.o) \bullet \Longrightarrow o} \text{ E_CBETA}$$

$$\frac{ealt_i = K\bar{x} \rightarrow o'}{\mathbf{case} \, K\bar{o} \, \mathbf{of} \, \overline{ealt} \Longrightarrow o'[\bar{o}/\bar{x}]} \text{ E_IOTA}$$

$$\frac{o_1 \Longrightarrow o'_1}{o_1 o_2 \Longrightarrow o'_1 o_2} \text{ E_APP_CONG}$$

$$\frac{o \Longrightarrow o'}{o \bullet \Longrightarrow o' \bullet} \text{ E_CAPP_CONG}$$

$$\frac{o \Longrightarrow o'}{\mathbf{case} \, o \, \mathbf{of} \, \overline{ealt} \Longrightarrow \mathbf{case} \, o' \, \mathbf{of} \, \overline{ealt}} \text{ E_CASE_CONG}$$

We translate from FC to the erased language via the erasure operation $o = |e|$:

$$\begin{aligned}
|x| &= x \\
|\lambda x:\tau.e| &= \lambda x.|e| \\
|\lambda c:\phi.e| &= \lambda \bullet.|e| \\
|e_1 e_2| &= |e_1| |e_2| \\
|\Lambda a:\kappa.e| &= |e| \\
|e \tau| &= |e| \\
|e \gamma| &= |e| \bullet \\
|K| &= K \\
|\mathbf{case}_\tau e \text{ of } \overline{alt}| &= \mathbf{case} |e| \text{ of } |\overline{alt}| \\
|e \triangleright \gamma| &= |e| \\
|\mathbf{contra} \gamma \tau| &= () \\
|K \bar{a} \bar{c} \bar{x} \rightarrow e| &= K \bar{x} \rightarrow |e|
\end{aligned}$$

In the **contra** case above, the right-hand side is the data constructor $()$ of the type `Unit`.

Lemma 62 (Erasing substitution). *For all e, τ , and a , $|e[\tau/a]| = |e|$.*

Lemma 63 (Erasing coercion substitution). *For all e, γ , and c , $|e[\gamma/c]| = |e|$.*

Lemma 64 (Erasing term substitution). *For all e, e' , and x , $|e[e'/x]| = |e|[[e'/x]]$.*

Lemma 65 (Erased values do not step). *If $o \Longrightarrow o'$, then o is not an erased value.*

Lemma 66 (Erased stepping is deterministic). *If $o \Longrightarrow o_1$ and $o \Longrightarrow o_2$, then $o_1 = o_2$.*

Theorem 67 (Type erasure). *If $e \xrightarrow{\bar{a}:\bar{\kappa}} e'$, then either $|e| \Longrightarrow |e'|$ or $|e| = |e'|$.*

Proof. By induction on $e \xrightarrow{\bar{a}:\bar{\kappa}} e'$.

Case S_BETA: We have $e = (\lambda x:\tau.e_1)e_2$ and $e' = e_1[e_2/x]$. Accordingly, we have $|e| = (\lambda x.|e_1|)|e_2|$ and $|e'| = |e_1[e_2/x]|$. By `E_BETA` and Lemma 64, we are done.

Case S_TBETA: The erasure of the term is unchanged.

Case S_CBETA: By `E_CBETA` and coercion substitution.

Case S_IOTA: The erased term steps by `E_IOTA`. We are done by use of the substitution lemmas above.

Case S_TRANS: The erasure of the term is unchanged.

Case S_TABS_CONG: By induction.

Case S_APP_CONG: By induction and `E_APP_CONG`.

Case S_TAPP_CONG: By induction.

Case S_CAPP_CONG: By induction and `E_CAPP_CONG`.

Case S_CASE_CONG: By induction and `E_CASE_CONG`.

Case S_CAST_CONG: By induction.

Case S_PUSH: The erasure of the term is unchanged.

Case S_TPUSH: The erasure of the term is unchanged.

Case S_CPUSH: The erasure of the term is unchanged.

Case S_APUSH: The erasure of the term is unchanged.

Case S_KPUSH: The erasure of the term is unchanged.

□

Lemma 68 (Erased redexes). *If $|e| = o$ and o is not an erased value, then e is not a value nor a coerced value.*

Proof. By induction on the structure of e .

Case $e = e_1 e_2$: For $|e_1 e_2| = o$ to be an erased value, $|e_1|$ must also be an erased value (headed by K). The induction hypothesis tells us that e_1 is not a value (and must not be headed by K). Thus, e is not a value.

Case $e = \Lambda a:\kappa.e_0$: Via the definition of erasure, we have $|e_0| = o$ and thus $|e_0|$ is not an erased value. The induction hypothesis tells us that e_0 must not then be a value, and thus e is not a value. (This depends on the fact that the definition of values requires a value in a type abstraction.)

Case $e = e_0 \tau$: Like the $e = e_1 e_2$ case.

Case $e = e_0 \gamma$: Like the $e = e_1 e_2$ case.

Case $e = e_0 \triangleright \gamma$: The induction hypothesis tells us that e_0 is not a value, and thus we are done.

Other cases: Trivial: either o is a value or e is not.

□

Theorem 69 (Types do not prevent evaluation). *If $\bar{a}:\bar{\kappa} \vdash e : \tau$ and $|e| \Longrightarrow o'$, then $e \xrightarrow{\bar{a}:\bar{\kappa}} e'$ and either $|e'| = o'$ or $|e'| = |e|$.*

Proof. Since $|e|$ steps, it must not be an erased value (Lemma 65). Thus, e must not be a value (Lemma 68). By the progress theorem (Theorem 51), we thus know that e must step to e' . By type erasure (Theorem 67), we can conclude that $|e'| = o'$ or $|e'| = |e|$, as desired. □