

Principles of Dependent Type Theory

Carlo Angiuli
Indiana University
cangiuli@iu.edu

Daniel Gratzer
Aarhus University
gratzer@cs.au.dk

(2024-08-30)

[Rake thudding against face]
Eeeuughhh

Robert Onderdonk Terwilliger Jr., Ph.D.
The Simpsons, Season 5 Episode 2, "Cape Feare."

Acknowledgements

We thank Lars Birkedal for his comments and suggestions on drafts of this book, and Sam Tobin-Hochstadt for many insightful conversations over lunch that helped us refine our narrative. We also thank the students who participated in *Modern Dependent Types* (CSCI-B619) at Indiana University and *Modern Dependent Type Theory* at Aarhus University in Spring 2024, for whom this book was prepared. A special thanks to Fred Fu, Rasmus Kirk Jakobsen, Artem Iurchenko, Pavel Kovalev, Kwing Hei Li, Amélia Liao, Mathias Adams Møller, Egor Namakonov, Thomas Porter, June Rousseau, Zixiu Su, Yafei Yang.

add names as people point out typos

Contents

Acknowledgements	iii
Contents	iv
Changes	vi
1 Introduction	1
1.1 Dependent types for functional programmers	3
2 Extensional type theory	16
2.1 The simply-typed lambda calculus	17
2.2 Towards the syntax of dependent type theory	24
2.3 The calculus of substitutions	28
2.4 Internalizing judgmental structure: $\Pi, \Sigma, \text{Eq}, \text{Unit}$	33
2.5 Inductive types: Void, Bool, Nat	44
2.6 Universes: U_0, U_1, U_2, \dots	55
2.7* Girard's paradox	63
2.8* Propositions and universes of propositions (DRAFT)	66
3 Metatheory and implementation	78
3.1 A judgmental reconstruction of proof assistants	79
3.2 Metatheory for type-checking	84
3.3* A case study in elaboration: definitions	91
3.4 Metatheory for computing	93
3.5* The set model of type theory	98
3.6 Equality in extensional type theory is undecidable	112
4 Intensional type theory	119
4.1 Programming with propositional equality	120
4.2 Intensional identity types	126
4.3 Limitations of the intensional identity type	132
4.4* Observational type theory (DRAFT)	143
5 Univalent type theories (DRAFT)	145
5.1 Propositions in intensional type theory	146
5.2 Homotopy type theory	150

5.3*	Motivating and justifying univalence	167
5.4	Cubical type theory	178
5.5*	Computing with coercions and compositions	195
A	Martin-Löf type theory	206
B	Solutions to selected exercises	215
	Bibliography	219

Changes

Below we summarize our major updates of this book.

2024-08-29

- Begin maintaining a change log.
- Added Section 3.5 on the set model of extensional type theory.
- Added solutions to selected exercises in Appendix B.
- Drafted Section 2.8 on propositions in extensional type theory.
- Drafted Chapter 5 on homotopy type theory.

2024-04-14

v0.1

- Added Chapter 1.
- Added Chapter 2 on extensional type theory.
- Added Chapter 3 on metatheory and implementation.
- Added Chapter 4 on intensional type theory.
- Added Appendix A collecting the formal rules of type theory.

Introduction

In this book, we aim to introduce the reader to a modern research perspective on the design of “full-spectrum” dependent type theories. After studying this book, readers should be prepared to engage with contemporary research papers on dependent type theory, and to understand the motivations behind recent extensions of Martin-Löf’s dependent type theory [ML84b], including observational type theory [AMS07], homotopy type theory [UF13], and cubical type theory [CCHM18; Ang+21].

This book is in an early draft form and is missing many relevant citations. The authors welcome any feedback.

Dependent type theory (henceforth just *type theory*) often appears arcane to outside observers for a handful of reasons. First, as in the parable of the elephant, there are myriad perspectives on type theory. The language presented in this book, *mutatis mutandis*, can be accurately described as:

- the core language of assertions and proofs in *proof assistants* like Agda [Agda], Coq [Coq], Lean [MU21], and Nuprl [Con+85];
- a richly-typed *functional programming language*, as in Idris [Bra13] and Pie [FC18], as well as in the aforementioned proof assistants Agda and Lean [Chr23].
- an *axiom system* for reasoning synthetically in a number of mathematical settings, including locally cartesian closed 1-categories [Hof95b], homotopy types [Shu21], and Grothendieck ∞ -topoi [Shu19];
- a structural [Tse17], constructive [ML82] *foundation for mathematics* as an alternative to ZFC set theory [Alt23].

A second difficulty is that it is quite complex to even *define* type theory in a precise fashion, for reasons we shall discuss in Section 2.2, and the relative merits of different styles of definition—and even which ones satisfactorily define any object whatsoever—have been the subject of great debate among experts over the years.

Finally, much of the literature on type theory is highly technical—involving either lengthy proofs by induction or advanced mathematical machinery—in order to account for its complex definition and applications. In this book we attempt to split the difference by presenting a mathematically-informed viewpoint on type theory while avoiding advanced mathematical prerequisites.

Goals of the course As researchers who work on designing new type theories, our goal in this course is to pose and begin to answer the following questions: *What makes a good type theory, and why are there so many?* We will focus on *notions of equality in Martin-Löf type theory* as a microcosm of this broader question, studying how extensional [ML82], intensional [ML75], observational [AMS07; SAG22; PT22], homotopy [UF13], and cubical type theories [CCHM18; Ang+21] have provided increasingly sophisticated answers to this deceptively simple question.

Although the design of type theory is inextricably linked to its applications (both theoretical and practical), we stress that this book focuses only on its design; there are many other resources for readers interested in learning how to use type theory.

Notes to the reader This book was written to accompany the authors' lectures in graduate courses on dependent type theory. As such, they are designed to be read in a linear fashion, with each chapter and section depending on many of the sections that come before it, with a few exceptions. Sections marked with *, such as Section 2.7, are considered optional and are not referenced later in the text; these sections cover topics that we consider important but nevertheless tangential to the main narrative. Smaller tangents are confined to *Remarks* and *Advanced Remarks*, the latter requiring more advanced mathematical prerequisites such as category theory. These often provide useful context or intuition but are again not integral to the main narrative.

Each chapter ends with a discussion of related literature, and we encourage the interested reader to follow these pointers to learn about these topics in greater depth. We have also attempted to include many references throughout the main body of the text, and the lengthy bibliography should also be considered a useful resource for further study.

Finally, we have included some exercises throughout the text to reinforce important concepts; for best results, the reader should work through at least some of these. Solutions to selected exercises can be found in Appendix B.

Notes to the expert We briefly remark on some editorial decisions that may surprise experts in type theory. First, we emphasize that this book is about the design of type theory, not how to use it. We therefore provide relatively few examples of working within type theory, focusing instead on type theories *qua* mathematical objects in their own right.

In light of this focus, experts may be surprised to find that our presentation does not explicitly rely on category theory. This was a difficult decision for the authors, both of whom view type theory from a categorical perspective, but we believe it is simply not feasible to insist that students begin their journey into type theory by first reading a book on category theory, and early attempts to simultaneously introduce category theory and type theory felt unsatisfactory on both counts.

That said, we do not attempt in any way to hide the presence of categories, functors, and naturality in the foundations of type theory. On the contrary, in Chapter 2 we define various connectives by the functors they (co)represent, phrased in more elementary language. We hope our exposition is accessible to readers encountering type theory for the first time, but also plainly categorical in flavor to those with more mathematical background.

Our perspective on type theory is deeply algebraic: we regard the judgments of type theory as being indexed by well-formed contexts and types, all defined only up to definitional equality. As a result, it is straightforward for us to introduce the notion of a model of type theory in Section 3.4, of which syntax is the initial example.

Finally, we have aimed to confine the non-optional sections of this book to fit within a semester of brisk lectures. For this reason we have elided numerous topics of interest, including a systematic treatment of inductive types, more discussion of elaboration, proofs of canonicity and normalization, and countless interesting variations of type theory.

In this chapter In Section 1.1 we introduce and motivate the concepts of type and term dependency, definitional equality, and propositional equality through the lens of typed functional programming. Note that Chapter 2 is self-contained albeit lacking in motivation, so readers unfamiliar with functional programming can safely skip ahead.

Goals of the chapter By the end of this chapter, you will be able to:

- Give examples of full-spectrum dependency.
- Explain the role of definitional equality in type-checking, and how and why it differs from ordinary closed-term evaluation.
- Explain the role of propositional equality in type-checking.

1.1 *Dependent types for functional programmers*

The reader is forewarned that the following section assumes some familiarity with functional programming, unlike the remainder of this book.

Types in programming languages For the purposes of this course, one should regard a programming language's (static) type system as its *grammar*, not as one of many potential static analyses that might be enabled or disabled.¹ Indeed, just as a parser may reject as nonsense a program whose parentheses are mismatched, or an untyped language's

¹The latter perspective is valid, but we wish to draw a sharp distinction between types *qua* (structural) grammar, and static analyses that may be non-local, non-structural, or non-substitutive in nature.

interpreter may reject as nonsense a program containing unbound identifiers, a type-checker may reject as nonsense the program `1 + "hi"` on the grounds that—much like the previous two examples—there is no way to successfully evaluate it.

Concretely, a type system divides a language’s well-parenthesized, well-scoped expressions into a collection of sets: the *expressions of type* `Nat` are those that “clearly” compute natural numbers, such as literal natural numbers (0, 1, 120), arithmetic expressions (`1 + 1`), and fully-applied functions that return natural numbers (`fact 5`, `atoi "120"`). Similarly, the expressions of type `String` are those that clearly compute strings (`"hi"`, `itoa 5`), and for any types `A` and `B`, the expressions of type `A → B` are those that clearly compute functions that, when passed an input of type `A`, clearly compute an output of type `B`.

What do we mean by “clearly”? One typically insists that type-checking be fully automated, much like parsing and identifier resolution. Given that determining the result of a program is in general undecidable, any automated type-checking process will necessarily compute a conservative underapproximation of the set of programs that compute (e.g.) natural numbers. (Likewise, languages may complain about unbound identifiers even in programs that can be evaluated without a runtime error!)

The goal of a type system is thus to rule out as many undesirable programs as possible without ruling out too many desirable ones, where both of these notions are subjective depending on which runtime errors one wants to rule out and which programming idioms one wants to support. Language designers engage in the neverending process of refining their type systems to rule out more errors and accept more correct code; full-spectrum dependent types can be seen as an extreme point in this design space.

1.1.1 Uniform dependency: length-indexed vectors

Every introduction to dependent types starts with the example of vectors, or lists with specified length. We start one step earlier by considering lists with a specified type of elements, a type which already exhibits a basic form of dependency.

Parameterizing by types One of the most basic data structures in functional programming languages is the *list*, which is either empty (written `[]`) or consists of an element `x` adjoined to a list `xs` (written `x :: xs`). In typed languages, we typically require that a list’s elements all have the same type so that we know what operations they support.

The simplest way to record this information is to have a separate type of lists for each type of element: a `ListOfNats` is either empty or a `Nat` and a `ListOfStrings`, a `ListOfStrings` is either empty or a `String` and a `ListOfStrings`, etc. This strategy clearly results in repetition at the level of the type system, but it also causes code duplication because operations that work uniformly for any type of elements (e.g., reversing a list) must be defined twice for the two apparently unrelated types `ListOfNats` and `ListOfStrings`.

In much the same way that functions—terms indexed by terms—promote code reuse by allowing programmers to write a series of operations once and perform them on many different inputs, we can solve both problems described above by allowing types and terms to be uniformly parameterized by types. Thus the types `ListOfNats` and `ListOfStrings` become two instances (`List Nat` and `List String`) of a single family of types `List`:²

```
data List (A : Set) : Set where
  [] : List A
  _::_ : A → List A → List A
```

and any operation that works for all element types A , such as returning the first (or all but first) element of a list, can be written as a family of operations:

```
head : (A : Set) → List A → A
head A [] = error "List must be non-empty."
head A (x :: xs) = x

tail : (A : Set) → List A → List A
tail A [] = error "List must be non-empty."
tail A (x :: xs) = xs
```

By partially applying `head` to its type argument, we see that `head Nat` has type `List Nat → Nat` and `head String` has type `List String → String`, and the expression `1 + (head Nat (1 :: []))` has type `Nat` whereas `1 + (head String ("hi" :: []))` is ill-typed because the second input to `+` has type `String`.

Parameterizing types by terms The perfectionist reader may find the `List A` type unsatisfactory because it does not prevent runtime errors caused by applying `head` and `tail` to the empty list `[]`. We cannot simply augment our types to track which lists are empty, because `2 :: 1 :: []` and `1 :: []` are both nonempty but we can apply `tail Nat` twice to the former before encountering an error, but only once to the latter.

Instead, we parameterize the type of lists not only by their type of elements as before but also by their length—a *term* of type `Nat`—producing the following family of types:³

```
data Vec (A : Set) : Nat → Set where
  [] : Vec A 0
  _::_ : {n : Nat} → A → Vec A n → Vec A (suc n)
```

Types parameterized by terms are known as *dependent types*.

²For the time being, the reader should understand $A : \text{Set}$ as notation meaning “ A is a type.”

³Curly braces $\{n : \text{Nat}\}$ indicate *implicit* arguments automatically inferred by the type-checker.

Now the types of concrete lists are more informative— $(2 :: 1 :: []) : \mathbf{Vec} \text{ Int } 2$ and $(1 :: []) : \mathbf{Vec} \text{ Int } 1$ —but more importantly, we can give head and tail more informative types which rule out the runtime error of applying them to empty lists. We do so by revising their input type to $\mathbf{Vec} A (\text{suc } n)$ for some $n : \mathbf{Nat}$, which is to say that the vector has length at least one, hence is nonempty:

```
head : {A : Set} {n : Nat} → Vec A (suc n) → A
-- head [] is impossible
head (x :: xs) = x
```

```
tail : {A : Set} {n : Nat} → Vec A (suc n) → Vec A n
-- tail [] is impossible
tail (x :: xs) = xs
```

Consider now the operation that concatenates two vectors:

```
append : {A : Set} {n : Nat} {m : Nat} → Vec A n → Vec A m → Vec A (n + m)
```

Unlike our previous examples, the output type of this function is indexed not by a variable A or n , nor a constant \mathbf{Nat} or 0 , nor even a constructor $\text{suc } -$, but by an *expression* $n + m$. This introduces a further complication, namely that we would like this expression to be simplified as soon as n and m are known. For example, if we apply `append` to two vectors of length one ($n = m = 1$), then the result will be a vector of length two ($n + m = 1 + 1 = 2$), and we would like the type system to be aware of this fact in the sense of accepting as well-typed the expression `head (tail (append l l'))` for l and l' of type $\mathbf{Vec} \text{ Nat } 1$.

Because `head (tail x)` is only well-typed when x has type $\mathbf{Vec} A (\text{suc } (\text{suc } n))$ for some $n : \mathbf{Nat}$, this condition amounts to requiring that the expression `append l l'` not only has type $\mathbf{Vec} A ((\text{suc } 0) + (\text{suc } 0))$ as implied by the type of `append`, but also type $\mathbf{Vec} A (\text{suc } (\text{suc } 0))$ as implied by its runtime behavior. In short, we would like the two type expressions $\mathbf{Vec} A (1 + 1)$ and $\mathbf{Vec} A 2$ to *denote the same type* by virtue of the fact that $1 + 1$ and 2 *denote the same value*. In practice, we achieve this by allowing the type-checker to *evaluate expressions in types during type-checking*.

In fact, the length of a vector can be any expression whatsoever of type \mathbf{Nat} . Consider `filter`, which takes a function $A \rightarrow \mathbf{Bool}$ and a list and returns the sublist for which the function returns true. If the input list has length n , what is the length of the output?

```
filter : {A : Set} {n : Nat} → (A → Bool) → Vec A n → Vec A ?
```

After a moment's thought we realize the length is not a function of n at all, but rather a recursive function of the input function and list:

```
filter : {A : Set} {n : Nat} → (f : A → Bool) → (l : Vec A n) → Vec A (filterLen f l)
```

```

filterLen : {A : Set} {n : Nat} → (A → Bool) → Vec A n → Nat
filterLen f [] = 0
filterLen f (x :: xs) = if f(x) then suc (filterLen f xs) else filterLen f xs

```

As before, once f and l are known the type of $\text{filter } f \ l : \text{Vec } A$ ($\text{filterLen } f \ l$) will simplify by evaluating $\text{filterLen } f \ l$, but as long as either remains a variable we cannot learn much by computation. Nevertheless, filterLen has many properties of interest: $\text{filterLen } f \ l$ is at most the length of l , $\text{filterLen } (\lambda x \rightarrow \text{false}) \ l$ is always 0 regardless of l , etc. We will revisit this point in Section 1.1.3.

Remark 1.1.1. If we regard Nat and $+$ as a user-defined data type and recursive function, as type theorists are wont to do, then filter 's type using filterLen is entirely analogous to append 's type using $+$. We wish to emphasize that, whereas one could easily imagine arithmetic being a privileged component of the type system, filter demonstrates that type indices may need to contain arbitrary user-defined recursive functions. \diamond

Another approach? If our only goal was to eliminate runtime errors from head and tail, we might reasonably feel that dependent types have overcomplicated the situation—we needed to introduce a new function just to write the type of filter ! And indeed there are simpler ways of keeping track of the length of lists, which we describe briefly here.

First let us observe that a lower bound on a list's length is sufficient to guarantee it is nonempty and thus that an application of head or tail will succeed; this allows us to trade precision for simplicity by restricting type indices to be arithmetic expressions. Secondly, in the above examples we can perform type-checking and “length-checking” in two separate phases, where the first phase replaces every occurrence of $\text{Vec } A \ n$ with $\text{List } A$ before applying a standard non-dependent type-checking algorithm. This is possible because we can regard the dependency in $\text{Vec } A \ n$ as expressing a computable *refinement*—or subset—of the non-dependent type of lists, namely $\{l : \text{List } A \mid \text{length } l = n\}$.

Combining these insights, we can by and large automate length-checking by recasting the type dependency of Vec in terms of arithmetic inequality constraints over an ML-style type system, and checking these constraints with SMT solvers and other external tools. At a very high level, this is the approach taken by systems such as Dependent ML [Xi07] and Liquid Haskell [Vaz+14]. Dependent ML, for instance, type-checks the usual definition of filter at the following type, without any auxiliary filterLen definition:

$$\text{filter} : \text{Vec } A \ m \rightarrow (\{n : \text{Nat} \mid n \leq m\} \times \text{Vec } A \ n)$$

Refinement type systems like these have proven very useful in practice and continue to be actively developed, but we will not discuss them any further for the simple reason

that, although they are a good solution to head/tail and many other examples, they cannot handle full-spectrum dependency as discussed below.

1.1.2 Non-uniform dependency: computing arities

Thus far, all our examples of (type- or term-) parameterized types are *uniformly* parameterized, in the sense that the functions $\text{List} : \text{Set} \rightarrow \text{Set}$ and $\text{Vec } A : \text{Nat} \rightarrow \text{Set}$ do not inspect their arguments; in contrast, ordinary term-level functions out of Nat such as $\text{fact} : \text{Nat} \rightarrow \text{Nat}$ can and usually do perform case-splits on their inputs. In particular, we have not yet considered any families of types in which the head, or top-level, type constructor (\rightarrow , Vec , Nat , etc.) differs between indices.

A type theory is said to have full-spectrum dependency if it permits the use of *non-uniformly term-indexed* families of types, such as the following Nat -indexed family:

```
nary : Set → Nat → Set
nary A 0 = A
nary A (suc n) = A → nary A n
```

Although Vec Nat and nary Nat are both functions $\text{Nat} \rightarrow \text{Set}$, the latter's head type constructor varies between indices: $\text{nary Nat } 0 = \text{Nat}$ but $\text{nary Nat } 1 = \text{Nat} \rightarrow \text{Nat}$.

Using nary to compute the type of n -ary functions, we can now define not only variadic functions but even higher-order functions taking variadic functions as input, such as apply which applies an n -ary function to a vector of length n :

```
apply : {A : Set} {n : Nat} → nary A n → Vec A n → A
apply x [] = x
apply f (x :: xs) = apply (f x) xs
```

For $A = \text{Nat}$ and $n = 1$, apply applies a unary function $\text{Nat} \rightarrow \text{Nat}$ to the head element of a $\text{Vec Nat } 1$; for $A = \text{Nat}$ and $n = 3$, it applies a ternary function $\text{Nat} \rightarrow \text{Nat} \rightarrow \text{Nat} \rightarrow \text{Nat}$ to the elements of a $\text{Vec Nat } 3$:

```
apply suc (1 :: []) : Nat -- evaluates to 2
apply _+_ : Vec Nat 2 → Nat
apply _+_ (1 :: 2 :: []) : Nat -- evaluates to 3
apply (λx y z → x + y + z) (1 :: 2 :: 3 :: []) : Nat -- evaluates to 6
```

Although apply is not the first time we have seen a function whose type involves a different recursive function—we saw this already with filter —this is our first example of a function that cannot be straightforwardly typed in an ML-style type system. Another way to put it is that $\text{nary } A n \rightarrow \text{Vec } A n \rightarrow A$ is not the refinement of an ML type because $\text{nary } A n$ is sometimes but not always a function type.

Remark 1.1.2. For the sake of completeness, it is also possible to consider *non-uniformly type-indexed* families of types, which go by a variety of names including non-parametric polymorphism, intensional type analysis, and typecase [HM95]. These often serve as optimized implementations of uniformly type-indexed families of types; a classic non-type-theoretic example is the C++ family of types `std::vector` for dynamically-sized arrays, whose `std::vector<bool>` instance may be compactly implemented using bitfields. \diamond

To understand the practical ramifications of non-uniform dependency, we will turn our attention to a more complex example: a basic implementation of `sprintf` in Agda (Figure 1.1). This function takes as input a **String** containing format specifiers such as `%u` (indicating a **Nat**) or `%s` (indicating a **String**), as well as additional arguments of the appropriate type for each format specifier present, and returns a **String** in which each format specifier has been replaced by the corresponding argument rendered as a **String**.

```
sprintf "%s %u" "hi" 2 : String -- evaluates to "hi 2"
sprintf "%s" : String → String
sprintf "nat %u then int %d then char %c" : Nat → Int → Char → String
sprintf "%u" 5 : String -- evaluates to "5"
sprintf "%u%% of %s%c" 3 "GD" 'P' : String -- evaluates to "3% of GDP"
```

Our implementation uses various types and functions imported from Agda's standard library, notably `toList : String → List Char` which converts a string to a list of characters (length-one strings `'x'`). It consists of four main components:

- a data type `Token` which enumerates all relevant components of the input **String**, namely format specifiers (such as `natTok : Token` for `%u` and `strTok : Token` for `%s`) and literal characters (`char 'x' : Token`);
- a function `lex` which tokenizes the input string, represented as a **List Char**, from left to right into a **List Token** for further processing;
- a function `args` which converts a **List Token** into a function type containing the additional arguments that `sprintf` must take; and
- the `sprintf` function itself.

Let us begin by convincing ourselves that our first example type-checks:

```
sprintf "%s %u" "hi" 2 : String -- evaluates to "hi 2"
```

Because `sprintf : (s : String) → printfType s`, the partial application `sprintf "%s %u"` has type `printfType "%s %u"`. By evaluation, the type-checker can see `printfType "%s %u" = args (strTok :: char ' ' :: natTok :: []) = String → Nat → String`. Thus `sprintf "%s %u" : String → Nat → String`, and the remainder of the expression type-checks easily.

```

data Token : Set where
  char : Char → Token
  intTok : Token
  natTok : Token
  chrTok : Token
  strTok : Token

lex : List Char → List Token
lex [] = []
lex ('%' :: '%' :: cs) = char '%' :: lex cs
lex ('%' :: 'd' :: cs) = intTok :: lex cs
lex ('%' :: 'u' :: cs) = natTok :: lex cs
lex ('%' :: 'c' :: cs) = chrTok :: lex cs
lex ('%' :: 's' :: cs) = strTok :: lex cs
lex (c :: cs) = char c :: lex cs

args : List Token → Set
args [] = String
args (char _ :: toks) = args toks
args (intTok :: toks) = Int → args toks
args (natTok :: toks) = Nat → args toks
args (chrTok :: toks) = Char → args toks
args (strTok :: toks) = String → args toks

printfType : String → Set
printfType s = args (lex (toList s))

sprintf : (s : String) → printfType s
sprintf s = loop (lex (toList s)) ""
  where
    loop : (toks : List Token) → String → args toks
    loop [] acc = acc
    loop (char c :: toks) acc = loop toks (acc ++ fromList (c :: []))
    loop (intTok :: toks) acc = λi → loop toks (acc ++ showInt i)
    loop (natTok :: toks) acc = λn → loop toks (acc ++ showNat n)
    loop (chrTok :: toks) acc = λc → loop toks (acc ++ fromList (c :: []))
    loop (strTok :: toks) acc = λs → loop toks (acc ++ s)

```

Figure 1.1: A basic Agda implementation of sprintf.

Now let us consider the definition of `sprintf`, which uses a helper function `loop` : $(\text{toks} : \text{List Token}) \rightarrow \text{String} \rightarrow \text{args toks}$ whose first argument stores the Tokens yet to be processed, and whose second argument is the **String** accumulated from printing the already-processed Tokens. What is needed to type-check the definition of `loop`? We can examine a representative case in which the next Token is `natTok`:

$$\text{loop } (\text{natTok} :: \text{toks}) \text{ acc} = \lambda n \rightarrow \text{loop } \text{toks} (\text{acc} ++ \text{showNat } n)$$

Note that `toks` : **List Token** and `acc` : **String** are (pattern) variables, and the right-hand side ought to have type `args (natTok :: toks)`. We can type-check the right-hand side—given that `_+_` : **String** \rightarrow **String** \rightarrow **String** is string concatenation and `showNat` : **Nat** \rightarrow **String** prints a natural number—and observe that it has type **Nat** \rightarrow `args toks` by the type of `loop`.

Type-checking this clause thus requires us to reconcile the right-hand side’s expected type `args (natTok :: toks)` with its actual type **Nat** \rightarrow `args toks`. Although these type expressions are quite dissimilar—one is a function type and the other is not—the definition of `args` contains a promising clause:

$$\text{args } (\text{natTok} :: \text{toks}) = \text{Nat} \rightarrow \text{args } \text{toks}$$

As in our earlier example of `Vec A (1+1)` and `Vec A 2` we would like the type expressions `args (natTok :: toks)` and **Nat** \rightarrow `args toks` to denote the same type, but unlike the equation $1 + 1 = 2$, here both sides contain a free variable `toks` so we cannot appeal to evaluation, which is a relation on *closed* terms (ones with no free variables).

One can nevertheless imagine some form of *symbolic evaluation* relation that extends evaluation to open terms and *can* equate these two expressions. In this particular case, this step of closed evaluation is syntactically indifferent to the value of `toks` and thus can be safely applied even when `toks` is a variable. (Likewise, to revisit an earlier example, the equation `filterLen f [] = 0` should hold even for variable `f`.)

Thus we would like the type expressions `args (natTok :: toks)` and **Nat** \rightarrow `args toks` to denote the same type by virtue of the fact that they *symbolically evaluate to the same symbolic value*, and to facilitate this we must allow the type-checker to *symbolically evaluate* expressions in types during type-checking. The congruence relation on expressions so induced is known as *definitional equality* because it contains defining clauses like this one.

Remark 1.1.3. Semantically we can justify this equation by observing that for any closed instantiation `toks` of `toks`, `args (natTok :: toks)` and **Nat** \rightarrow `args toks` will evaluate to the same type expression—at least, once we have defined evaluation of type expressions—and thus this equation always holds at runtime. But just as (for reasons of decidability) the condition “when this expression is applied to a natural number it evaluates to a natural number” is a necessary but not sufficient condition for type-checking at **Nat** \rightarrow **Nat**, we do not want to take this semantic condition as the definition of definitional equality. It

is however a necessary condition assuming that the type system is sound for the given evaluation semantics. (See Section 3.4.) \diamond

Definitional equality is the central concept in full-spectrum dependent type theory because it determines which types are equal and thus which terms have which types. In practice, it is typically defined as the congruence closure of the β -like reductions (also known as $\beta\delta\zeta\iota$ -reductions) plus η -equivalence at some types; see Chapter 2 for details.

1.1.3 Proving type equations

Unfortunately, in light of Remark 1.1.3, there are many examples of type equations that are not direct consequences of ordinary or even symbolic evaluation. On occasion these equations are of such importance that researchers may attempt to make them definitional—that is, to include them in the definitional equality relation and adjust the type-checking algorithm accordingly [AMB13]. But such projects are often major research undertakings, and there are even examples of equations that can be definitional but are in practice best omitted due to efficiency or usability issues [Alt+01].

Let us turn once again to the example of `filter` from Section 1.1.2.

$$\text{filter} : \{A : \text{Set}\} \{n : \text{Nat}\} \rightarrow (f : A \rightarrow \text{Bool}) \rightarrow (l : \text{Vec } A \ n) \rightarrow \text{Vec } A \ (\text{filterLen } f \ l)$$

$$\text{filterLen} : \{A : \text{Set}\} \{n : \text{Nat}\} \rightarrow (A \rightarrow \text{Bool}) \rightarrow \text{Vec } A \ n \rightarrow \text{Nat}$$

$$\text{filterLen } f \ [] = 0$$

$$\text{filterLen } f \ (x :: xs) = \text{if } f(x) \ \text{then } \text{suc } (\text{filterLen } f \ xs) \ \text{else } \text{filterLen } f \ xs$$

Suppose for the sake of argument that we want the operation of filtering an arbitrary vector by the constantly false predicate to return a `Vec A 0`:

$$\text{filterAll} : \{A : \text{Set}\} \{n : \text{Nat}\} \rightarrow \text{Vec } A \ n \rightarrow \text{Vec } A \ 0$$

$$\text{filterAll } l = \text{filter } (\lambda x \rightarrow \text{false}) \ l \quad \text{-- does not type-check}$$

The right-hand side above has type `Vec A (filterLen (λx → false) l)` rather than `Vec A 0` as desired, and in this case the expression `filterLen (λx → false) l` cannot be simplified by (symbolic) evaluation because `filterLen` computes by recursion on `l` which is a variable. However, by induction on the possible instantiations of `l : Vec A n`, either:

- `l = []`, in which case `filterLen (λx → false) []` is definitionally equal (in fact, evaluates) to 0; or

- $l = x :: xs$, in which case we have the definitional equalities

$$\begin{aligned} & \text{filterLen } (\lambda x \rightarrow \text{false}) (x :: xs) \\ &= \text{if false then suc (filterLen } (\lambda x \rightarrow \text{false}) xs) \text{ else filterLen } (\lambda x \rightarrow \text{false}) xs \\ &= \text{filterLen } (\lambda x \rightarrow \text{false}) xs \end{aligned}$$

for any x and xs . By the inductive hypothesis on xs , $\text{filterLen } (\lambda x \rightarrow \text{false}) xs = 0$ and thus $\text{filterLen } (\lambda x \rightarrow \text{false}) (x :: xs) = 0$ as well.

By adding a type of *provable equations* $a \equiv b$ to our language, we can compactly encode this inductive proof as a recursive function computing $\text{filterLen } (\lambda x \rightarrow \text{false}) l \equiv 0$:

```
_≡_ : {A : Set} → A → A → Set
refl : {A : Set} {x : A} → x ≡ x

lemma : {A : Set} {n : Nat} → (l : Vec A n) → filterLen (λl → false) l ≡ 0
lemma [] = refl
lemma (x :: xs) = lemma xs
```

The `[]` clause of `lemma` ought to have type $\text{filterLen } (\lambda l \rightarrow \text{false}) [] \equiv 0$, which is definitionally equal to the type $0 \equiv 0$ and thus `refl` type-checks. The $(x :: xs)$ clause must have type $\text{filterLen } (\lambda l \rightarrow \text{false}) (x :: xs) \equiv 0$, which is definitionally equal to $\text{filterLen } (\lambda l \rightarrow \text{false}) xs \equiv 0$, the expected type of the recursive call `lemma xs`.

Now armed with a function `lemma` that constructs for any $l : \text{Vec } A \ n$ a proof that $\text{filterLen } (\lambda l \rightarrow \text{false}) l \equiv 0$, we can justify *casting* from the type $\text{Vec } A \ (\text{filterLen } (\lambda l \rightarrow \text{false}) l)$ to $\text{Vec } A \ 0$. The dependent casting operation that passes between provably equal indices of a dependent type (in this case $\text{Vec } A : \text{Nat} \rightarrow \text{Set}$) is typically called **subst**:

```
subst : {A : Set} {x y : A} → (P : A → Set) → x ≡ y → P(x) → P(y)

filterAll : {A : Set} {n : Nat} → Vec A n → Vec A 0
filterAll {A} l = subst (Vec A) (lemma l) (filter (λx → false) l)
```

Remark 1.1.4. The **subst** operation above is a special case of a much stronger principle stating that the two types $P(x)$ and $P(y)$ are *isomorphic* whenever $x \equiv y$: we can not only cast $P(x) \rightarrow P(y)$ but also $P(y) \rightarrow P(x)$ by symmetry of equality, and both round trips cancel. So although a proof $x \equiv y$ does not make $P(x)$ and $P(y)$ definitionally equal, they are nevertheless equal in the sense of having the same elements up to isomorphism. \diamond

Uses of **subst** are very common in dependent type theory; because dependently-typed functions can both require and ensure complex invariants, one must frequently prove that

the output of some function is a valid input to another.⁴ Crucially, although **subst** is an “escape hatch” that compensates for the shortcomings of definitional equality, it cannot result in runtime errors—unlike explicit casts in most programming languages—because casting from $P(x)$ to $P(y)$ requires a machine-checked proof that $x \equiv y$. We can ask for such proofs because dependent type theory is not only a functional programming language but also a higher-order intuitionistic logic that can express inductive proofs of type equality, and as we saw with `filterAll`, its type-checker serves also as a proof-checker.

The dependent type $x \equiv y$ is known as *propositional equality*, and it is perhaps the second most important concept in dependent type theory because it is the source of all non-definitional type equations visible within the theory. There are many formulations of propositional equality; they all implement `_≡_`, `refl`, and `subst` but differ in many other respects, and each has unique benefits and drawbacks. We will discuss propositional equality at length in Chapters 4 and 5.

To foreshadow the design space of propositional equality, consider that the `subst` operator may itself be subject to various definitional equalities. If we apply `filterAll` to a closed list `ls`, then lemma `ls` will evaluate to `refl`, so `filterAll ls` is definitionally equal to `subst (Vec A) refl (filter (λx → false) ls)`. At this point, `filter (λx → false) ls` already has the desired type `Vec A 0` because `filterLen (λx → false) ls` evaluates to 0, and thus the two types involved in the cast are now definitionally equal. Ideally the `subst` term would now disappear having completed its job, and indeed the corresponding definitional equality `subst P refl x = x` does hold for many versions of propositional equality.

⁴A more realistic variant of our lemma might account for any predicate that returns false on all the elements of the given list, not just the constantly false predicate. Alternatively, one might prove that for any `s : String`, the final return type of `sprintf s` is `String`.

Further reading

Our four categories of dependency—types/terms depending on types/terms—are reminiscent of the λ -cube of generalized type systems in which one augments the simply-typed λ -calculus (whose functions exhibit term-on-term dependency) with any combination of the remaining three forms of dependency [Bar91]; adding all three yields the full-spectrum dependent type theory known as the calculus of constructions [CH88]. However, the technical details of this line of work differ significantly from our presentation in Chapter 2.

The remarkable fact that type theory is both a functional programming language and a logic is known by many names including *the Curry–Howard correspondence* and *propositions as types*. It is a very broad topic with many treatments; book-length expositions include *Proofs and Types* [GLT89] and *PROGRAM = PROOF* [Mim20].

The code in this chapter is written in Agda syntax [Agda]. For more on dependently-typed programming in Agda, see *Verified Functional Programming in Agda* [Stu16]; for a more engineering-oriented perspective on dependent types, see *Type-Driven Development with Idris* [Bra17]. The `sprintf` example in Section 1.1.2 is inspired by the paper *Cayenne — A Language with Dependent Types* [Aug99]. Conversely, to learn about using Agda as a proof assistant for programming language theory, see *Programming Language Foundations in Agda* [WKS22].

Extensional type theory

In order to understand the subtle differences between modern dependent type theories, we must first study the formal definition of a dependent type theory as a mathematical object. We will then be prepared for Chapter 3, in which we study mathematical properties of type theory—and particularly of definitional and propositional equality—and their connection to computer implementations of type theory. In this chapter we therefore present the judgmental theory of Martin-Löf’s *extensional type theory* [ML82], one of the canonical variants of dependent type theory. We strongly suggest following the exposition rather than simply reading the rules, but the rules are collected for convenience in Appendix A (ignoring the rules marked with (TT), which are present only in intensional type theory).

Given the time constraints of this course, we do not attempt to give a comprehensive account of the syntax of type theories, nor do we present any of the many alternative methods of defining type theory, some of which are more efficient (but more technical) than the one we present here. These questions lead to the fascinating and deep area of *logical frameworks* which we must regrettably leave for a different course.

In this chapter In Section 2.1 we recall the concepts of judgments and inference rules in the setting of the simply-typed lambda calculus. In Section 2.2 we consider how to adapt these methods to the dependent setting, and in Section 2.3 we develop these ideas into the basic judgmental structure of dependent type theory, in which substitution plays a key role. In Section 2.4 we extend the basic rules of type theory with rules governing dependent products, dependent sums, extensional equality, and unit types. We argue that these connectives can be understood as *internalizations of judgmental structure*, a perspective which provides a conceptual justification of these connectives’ rules. In Section 2.5 we define several inductive types—the empty type, booleans, and natural numbers—and explain how and why these types do not fit the pattern of the previous section. Finally, in Section 2.6 we discuss large elimination, which is implicit in our examples of full-spectrum dependency from Section 1.1, and its internalization via universe types.

add descriptions of new sections, and explain optional sections

Goals of the chapter By the end of this chapter, you will be able to:

- Define the core judgments of dependent type theory, and explain how and why they differ from the judgments of simple type theory.
- Explain the role of substitutions in the syntax of dependent type theory.

- Define and justify the rules of the core connectives of type theory.

2.1 The simply-typed lambda calculus

The theory of typed functional programming is built on extensions of a core language known as the *simply-typed lambda calculus*, which supports two types of data:

- functions of type $A \rightarrow B$ (for any types A, B): we write $\lambda x.b$ for the function that sends any input x of type A to an output b of type B , and write $f a$ for the application of a function f of type $A \rightarrow B$ to an input a of type A ; and
- ordered pairs of type $A \times B$ (for any types A, B): we write (a, b) for the pair of a term a of type A with a term b of type B , and write $\text{fst}(p)$ and $\text{snd}(p)$ respectively for the first and second projections of a pair p of type $A \times B$.

It can also be seen as the implication–conjunction fragment of intuitionistic propositional logic, or as an axiom system for cartesian closed categories.

In this section we formally define the simply-typed lambda calculus as a collection of judgments presented by inference rules, in order to prepare ourselves for the analogous—but considerably more complex—definition of dependent type theory in the remainder of this chapter. Our goal is thus not to give a textbook account of the simply-typed lambda calculus but to draw the reader’s attention to issues that will arise in the dependent setting.

Readers familiar with the simply-typed lambda calculus should be aware that our definition does not reference the untyped lambda calculus (as discussed in Remark 2.1.2) and considers terms modulo $\beta\eta$ -equivalence (Section 2.1.2).

2.1.1 Contexts, types, and terms

The simply-typed lambda calculus is made up of two *sorts*, or grammatical categories, namely types and terms. We present these sorts by two well-formedness *judgments*:

- the judgment $A \text{ type}$ stating that A is a well-formed type, and
- for any well-formed type A , the judgment $a : A$ stating that a is a well-formed term of that type.

By comprehension these judgments determine respectively the collection of well-formed types and, for every element of that collection, the collection of well-formed terms of that type. (From now on we will stop writing “well-formed” because we do not consider any other kind of types or terms; see Remark 2.1.2.)

Remark 2.1.1. A judgment is simply a proposition in our ambient mathematics, one which takes part in the definition of a logical theory; we use this terminology to distinguish such meta-propositions from the propositions of the logic that is being defined [ML87]. Similarly, a sort is a type in the ambient mathematics, as distinguished from the types of the theory being defined. We refer to the ambient mathematics (in which our definition is being carried out) as the *metatheory* and the logic being defined as the *object theory*.

In this course we will be relatively agnostic about our metatheory, which the reader can imagine as “ordinary mathematics.” However, one can often simplify matters by adopting a domain-specific metatheory (a *logical framework*) well-suited to defining languages/logics, as an additional level of indirection within the ambient metatheory. \diamond

Types We can easily define the types as the expressions generated by the following context-free grammar:

$$\text{Types } A, B := \mathbf{b} \mid A \times B \mid A \rightarrow B$$

We say that the judgment A type (“ A is a type”) holds when A is a type in the above sense. Note that in addition to function and product types we have included a base type \mathbf{b} ; without \mathbf{b} the grammar would have no terminal symbols and would thus be empty.

Equivalently, we could define the A type judgment by three *inference rules* corresponding to the three production rules in the grammar of types:

$$\frac{}{\mathbf{b} \text{ type}} \qquad \frac{A \text{ type} \quad B \text{ type}}{A \times B \text{ type}} \qquad \frac{A \text{ type} \quad B \text{ type}}{A \rightarrow B \text{ type}}$$

Each inference rule has some number of premises (here, zero or two) above the line and a single conclusion below the line; by combining these rules into trees whose leaves all have no premises, we can produce *derivations* of judgments (here, the well-formedness of a type) at the root of the tree. The tree below is a proof that $(\mathbf{b} \times \mathbf{b}) \rightarrow \mathbf{b}$ is a type:

$$\frac{\frac{\frac{}{\mathbf{b} \text{ type}} \quad \frac{}{\mathbf{b} \text{ type}}}{\mathbf{b} \times \mathbf{b} \text{ type}} \quad \frac{}{\mathbf{b} \text{ type}}}{(\mathbf{b} \times \mathbf{b}) \rightarrow \mathbf{b} \text{ type}}}$$

Terms Terms are considerably more complex than types, so before attempting a formal definition we will briefly summarize our intentions. For the remainder of this section, fix a finite set I . The well-formed terms are as follows:

- for any $i \in I$, the base term \mathbf{c}_i has type \mathbf{b} ;

- pairing (a, b) has type $A \times B$ when $a : A$ and $b : B$;
- first projection $\text{fst}(p)$ has type A when $p : A \times B$;
- second projection $\text{snd}(p)$ has type B when $p : A \times B$;
- a function $\lambda x.b$ has type $A \rightarrow B$ when $b : B$ where b can contain (in addition to the usual term formers) the variable term $x : A$ standing for the function's input; and
- a function application $f a$ has type B when $f : A \rightarrow B$ and $a : A$.

The first difficulty we encounter is that unlike types, which are a single sort, there are infinitely many sorts of terms (one for each type) many of which refer to one another. A more significant issue is to make sense of the clause for functions: the body b of a function $\lambda x.b : A \rightarrow B$ is a term of type B according to our original grammar *extended by* a new constant $x : A$ representing an indeterminate term of type A . Because b can again be or contain a function $\lambda y.c$, we must account for finitely many extensions $x : A, y : B, \dots$

To account for these extensions we introduce an auxiliary sort of *contexts*, or lists of variables paired with types, representing local extensions of our theory by variable terms.

Contexts The judgment $\vdash \Gamma \text{ cx}$ (“ Γ is a context”) expresses that Γ is a list of pairs of term variables with types. We write $\mathbf{1}$ for the empty context and $\Gamma, x : A$ for the extension of Γ by a term variable x of type A . As a context-free grammar, we might write:

$$\begin{array}{l} \text{Variables } x, y := x \mid y \mid z \mid \dots \\ \text{Contexts } \Gamma := \mathbf{1} \mid \Gamma, x : A \end{array}$$

Equivalently, in inference rule notation:

$$\frac{}{\vdash \mathbf{1} \text{ cx}} \qquad \frac{\vdash \Gamma \text{ cx} \quad A \text{ type}}{\vdash \Gamma, x : A \text{ cx}}$$

We will not spend time discussing variables or binding in this book because variables will, perhaps surprisingly, not be a part of our definition of dependent type theory. For the purposes of this section we will simply assume that there is an infinite set of variables x, y, z, \dots , and that all the variables in any given context or term are distinct.

Terms revisited With contexts in hand we are now ready to define the term judgment, which we revise to be relative to a context Γ . The judgment $\Gamma \vdash a : A$ (“ a has type A in context Γ ”) is defined by the following inference rules:

$$\begin{array}{c}
\frac{(x : A) \in \Gamma}{\Gamma \vdash x : A} \quad \frac{i \in I}{\Gamma \vdash \mathbf{c}_i : \mathbf{b}} \quad \frac{\Gamma \vdash a : A \quad \Gamma \vdash b : B}{\Gamma \vdash (a, b) : A \times B} \quad \frac{\Gamma \vdash p : A \times B}{\Gamma \vdash \mathbf{fst}(p) : A} \\
\\
\frac{\Gamma \vdash p : A \times B}{\Gamma \vdash \mathbf{snd}(p) : B} \quad \frac{\Gamma, x : A \vdash b : B}{\Gamma \vdash \lambda x. b : A \rightarrow B} \quad \frac{\Gamma \vdash f : A \rightarrow B \quad \Gamma \vdash a : A}{\Gamma \vdash f a : B}
\end{array}$$

The rules for \mathbf{c}_i , pairing, projections, and application straightforwardly render our text into inference rule form, framed by a context Γ that is unchanged from premises to conclusion. The lambda rule explains how contexts are changed: the body of a lambda is typed in an extended context; and the variable rule explains how contexts are used: in context Γ , the variables of type A in Γ serve as additional terminal symbols of type A .

Rules such as pairing or lambda that describe how to create terms of a given type former are known as *introduction* rules, and rules describing how to use terms of a given type former, like projection and application, are known as *elimination* rules.

Remark 2.1.2. An alternative approach that is perhaps more familiar to programming languages researchers is to define a collection of *preterms*

$$\text{Terms } a, b := \mathbf{c}_i \mid x \mid (a, b) \mid \mathbf{fst}(a) \mid \mathbf{snd}(a) \mid \lambda x. a \mid a b$$

which includes ill-formed (typeless) terms like $\mathbf{fst}(\lambda x. x)$ in addition to the well-formed (typed) ones captured by our grammar above, and the inference rules are regarded as carving out various subsets of well-formed terms [Har16]. In fact, one often gives computational meaning to *all* preterms (as an extension of the untyped lambda calculus) and then proves that the well-typed ones are in some sense computationally well-behaved.

This is *not* the approach we are taking here; to us the term expression $\mathbf{fst}(\lambda x. x)$ does not exist any more than the type expression $\rightarrow \times \rightarrow$.¹ In fact, in light of Section 2.1.2, there will not even exist a “forgetful” map from our collections of terms to these preterms. \diamond

2.1.2 Equational rules

One shortcoming of our definition thus far is that our projections don’t actually project anything and our function applications don’t actually apply functions—there is no sense yet in which $\mathbf{fst}((a, b)) : A$ or $(\lambda x. x) a : A$ “are” $a : A$. Rather than equip our terms with operational meaning, we will *quotient* our terms by equations that capture a notion of sameness including these examples. The reader can imagine this process as analogous to

¹Perhaps one’s definition of context-free grammar carves out the grammatical expressions out of arbitrary strings over an alphabet, but this process occurs at a different level of abstraction. The reader should banish such thoughts along with their thoughts about terms with mismatched parentheses.

the presentation of algebras by *generators and relations*, in which our terms thus far are the generators of a “free algebra” of (well-formed but) uninterpreted expressions.

Our true motivation for this quotient is to anticipate the definitional equality of dependent type theory, but there are certainly intrinsic reasons as well, perhaps most notably that the quotiented terms of the simply-typed lambda calculus serve as an axiom system for reasoning about cartesian closed categories [Cro94, Chapter 4].

We quotient by the congruence relation generated by the following rules:

$$\frac{\Gamma \vdash a : A \quad \Gamma \vdash b : B}{\Gamma \vdash \mathbf{fst}((a, b)) = a : A} \quad \frac{\Gamma \vdash a : A \quad \Gamma \vdash b : B}{\Gamma \vdash \mathbf{snd}((a, b)) = b : B} \quad \frac{\Gamma \vdash p : A \times B}{\Gamma \vdash p = (\mathbf{fst}(p), \mathbf{snd}(p)) : A \times B}$$

$$\frac{\Gamma, x : A \vdash b : B \quad \Gamma \vdash a : A}{\Gamma \vdash (\lambda x. b) a = b[a/x] : B} \quad \frac{\Gamma \vdash f : A \rightarrow B}{\Gamma \vdash f = \lambda x. (f x) : A \rightarrow B}$$

The equations pertaining to elimination after introduction (projection from pairs and application of lambdas) are called *β -equivalences*; the equations pertaining to introduction after elimination (pairs of projections and lambdas of applications) are *η -equivalences*.

We emphasize that these equations are not *a priori* directed, and are not restricted to the “top level” of terms; we genuinely take the quotient of the collection of terms at each type by these equations, automatically inducing equations such as $\lambda x. x = \lambda x. \mathbf{fst}((x, x))$.

The first two rules explain that projecting from a pair has the evident effect. The third rule states that every term of type $A \times B$ can be written as a pair (of its projections), in effect transforming the introduction rule for products from merely a sufficient condition to a necessary one as well. Similarly, the fifth rule states that every $f : A \rightarrow B$ can be written as a lambda (of its application).

The fourth rule explains that applying a lambda function $\lambda x. b$ to an argument a is equal to the body b of that lambda with all occurrences of the placeholder variable x replaced by the term a . However, this equation makes reference to a *substitution* operation $b[a/x]$ (“substitute a for x in b ”) that we have not yet defined.

Substitution We can define substitution $b[a/x]$ by structural recursion on b :

$$\begin{aligned}
\mathbf{c}_i[c/x] &:= \mathbf{c}_i \\
x[c/x] &:= c \\
y[c/x] &:= y && \text{(for } x \neq y\text{)} \\
(a, b)[c/x] &:= (a[c/x], b[c/x]) \\
\mathbf{fst}(p)[c/x] &:= \mathbf{fst}(p[c/x]) \\
\mathbf{snd}(p)[c/x] &:= \mathbf{snd}(p[c/x]) \\
(\lambda y. b)[c/x] &:= \lambda y. b[c/x] && \text{(for } x \neq y\text{)} \\
(f a)[c/x] &:= f[c/x] a[c/x]
\end{aligned}$$

In the case of substituting into a lambda $(\lambda y. b)[c/x]$, we assume that the bound variable y introduced by the lambda is different from the variable x being substituted away. In practice they may coincide, in which case one must rename y (and all references to y in b) before applying this rule. In any case, we intend this substitution to be *capture-avoiding* in the sense of not inadvertently changing the referent of bound variables.

However, because we have quotiented our collection of terms by $\beta\eta$ -equivalence, it is not obvious that substitution is well-defined as a function out of the collection of terms; in order to map out of the quotient, we must check that substitution behaves equally on equal terms. (It is also not obvious that substitution is a function *into* the collection of terms, in the sense of producing well-formed terms, as we will discuss shortly.)

Consider the equation $\mathbf{fst}((a, b)) = a$. To see that substitution respects this equation, we can substitute into the left-hand side, yielding:

$$(\mathbf{fst}((a, b)))[c/x] = \mathbf{fst}((a, b)[c/x]) = \mathbf{fst}((a[c/x], b[c/x]))$$

which is β -equivalent to $a[c/x]$, the result of substituting into the right-hand side. We can check the remaining equations in a similar fashion; the $x \neq y$ condition on substitution into lambdas is necessary for substitution to respect β -equivalence of functions.

2.1.3 Who type-checks the typing rules?

Our stated goal in Section 2.1.1 was to define a collection of well-formed types (written A type), and for each of these a collection of well-formed terms (written $a : A$). Have we succeeded? First of all, our definition of terms is now indexed by contexts Γ and written $\Gamma \vdash a : A$, to account for variables introduced by lambdas. This is no problem: we recover the original notion of (closed) term by considering the empty context $\mathbf{1}$. Nor is there any issue defining the collections of types $\mathbf{Ty} = \{A \mid A \text{ type}\}$ and contexts $\mathbf{Cx} = \{\Gamma \mid \vdash \Gamma \text{ cx}\}$ as presented by the grammars or inference rules in Section 2.1.1.

It is less clear that the collections of *terms* are well-defined. We would like to say that the collection of terms of type A in context Γ , $\text{Tm}(\Gamma, A)$, is the set of a for which there exists a derivation of $\Gamma \vdash a : A$, modulo the relation $a \sim b \iff$ there exists a derivation of $\Gamma \vdash a = b : A$. Several questions arise immediately; for instance, is it the case that whenever $\Gamma \vdash a : A$ is derivable, Γ is a context and A is a type? If not, then we have some “junk” judgments that should not correspond to elements of some $\text{Tm}(\Gamma, A)$.

Lemma 2.1.3. *If $\Gamma \vdash a : A$ then $\vdash \Gamma$ cx and A type.*

To prove such a statement, one proceeds by induction on derivations of $\Gamma \vdash a : A$. If, say, the derivation ends as follows:

$$\frac{\vdots}{\Gamma \vdash p : A \times B} \\ \Gamma \vdash \text{fst}(p) : A$$

then the inductive hypothesis applied to the derivation of $\Gamma \vdash p : A \times B$ tells us that $\vdash \Gamma$ cx and $A \times B$ type. The former is exactly one of the two statements we are trying to prove. The other, A type, follows from an “inversion lemma” (proven by cases on the $-$ type judgment) that A type is not only a sufficient but also a necessary condition for $A \times B$ type.

Unfortunately our proof runs into an issue at the base cases, or at least it is not clear over what Γ the following rules range:

$$\frac{(x : A) \in \Gamma}{\Gamma \vdash x : A} \qquad \frac{i \in I}{\Gamma \vdash \mathbf{c}_i : \mathbf{b}}$$

We must either add premises to these rules stating $\vdash \Gamma$ cx, or else clarify that Γ always ranges only over contexts (which will be our strategy moving forward; see Notation 2.2.1).

Another question is the well-definedness of our quotient:

Lemma 2.1.4. *If $\Gamma \vdash a = b : A$ then $\Gamma \vdash a : A$ and $\Gamma \vdash b : A$.*

But because β -equivalence refers to substitution, proving this lemma requires:

Lemma 2.1.5 (Substitution). *If $\Gamma, x : A \vdash b : B$ and $\Gamma \vdash a : A$ then $\Gamma \vdash b[a/x] : B$.*

We already saw that we must check that substitution $b[a/x]$ respects equality of b , but we must also check that it produces well-formed terms, again by induction on b . Note that substitution changes a term’s context because it eliminates one of its free variables.

If we resume our attempt to prove Lemma 2.1.4, we will notice that substitution is not the only time that the context of a term changes; in the right-hand side of the η -rule of functions, f is in context $\Gamma, x : A$, whereas in the premise and left-hand side it is in Γ :

$$\frac{\Gamma \vdash f : A \rightarrow B}{\Gamma \vdash f = \lambda x. (f x) : A \rightarrow B}$$

And thus we need yet another lemma.

Lemma 2.1.6 (Weakening). *If $\Gamma \vdash b : B$ and $\Gamma \vdash A$ type then $\Gamma, x : A \vdash b : B$.*

We will not belabor the point any further; eventually one proves enough lemmas to conclude that we have a set of contexts Cx , a set of types Ty , and for every $\Gamma \in Cx$ and $A \in Ty$ a set of terms $Tm(\Gamma, A)$. The complexity of each result is proportional to the complexity of that sort's definition: we define types outright, contexts by simple reference to types, and terms by more complex reference to both types and contexts. The judgments of dependent type theory are both more complex and more intertwined; rather than enduring proportionally more suffering, we will adopt a slightly different approach.

Finally, whereas all the metatheorems mentioned in this section serve only to establish that our definition is mathematically sensible, there are more genuinely interesting and contentful metatheorems one might wish to prove, including *canonicity*, the statement that (up to equality) the only closed terms of \mathbf{b} are of the form \mathbf{c}_i (i.e., $Tm(\mathbf{1}, \mathbf{b}) = \{\mathbf{c}_i\}_{i \in I}$), and *decidability of equality*, the statement that for any $\Gamma \vdash a : A$ and $\Gamma \vdash b : A$ we can write a program which determines whether or not $\Gamma \vdash a = b : A$.

2.2 Towards the syntax of dependent type theory

The reader is forewarned that the rules in this section serve to bridge the gap between Section 2.1 and our “official” rules for extensional type theory, which start in Section 2.3.

As we discussed in Section 1.1, the defining distinction between dependent and simple type theory is that in the former, types can contain term expressions and even term variables. Thus, whereas in Section 2.1 a simple context-free grammar sufficed to define the collection of types and we needed a context-sensitive system of inference rules to define the well-typed terms, in dependent type theory we will find that both the types and terms are context-sensitive because they refer to one another.

Types and contexts When is the dependent function type $(x : A) \rightarrow B$ well-formed? Certainly A and B must be well-formed types, but B is allowed to contain the term variable $x : A$ whereas A is not. In the case of $(n : \text{Nat}) \rightarrow \text{Vec String}$ ($\text{suc } n$), the well-formedness of the codomain depends on the fact that $\text{suc } n$ is a well-formed term of type Nat (the

indexing type of `Vec String`), which in turn depends on the fact that n is known to be an expression (in particular, a variable) of type `Nat`.

Thus as with the *term* judgment of Section 2.1, the *type* judgment of dependent type theory must have access to the context of term variables, so we replace the A type judgment (“ A is a type”) of the simply-typed lambda calculus with a judgment $\Gamma \vdash A \text{ type}$ (“ A is a type in context Γ ”). This innocuous change has many downstream implications, so we will be fastidious about the context in which a type is well-formed.

The first consequence of this change is that contexts of term variables, which we previously defined simply as lists of well-formed types, must now also take into account *in what context* each type is well-formed. Informally we say that each type can depend on all the variables before it in the context; formally, one might define the judgment $\vdash \Gamma \text{ cx}$ by the following pair of rules:

$$\frac{}{\vdash \mathbf{1} \text{ cx}} \qquad \frac{\vdash \Gamma \text{ cx} \quad \Gamma \vdash A \text{ type}}{\vdash \Gamma, x : A \text{ cx}}$$

Notice that the rules defining the judgment $\vdash \Gamma \text{ cx}$ refer to the judgment $\Gamma \vdash A \text{ type}$, which in turn depends on our notion of context. This kind of mutual dependence will continue to crop up throughout the rules of dependent type theory.

Notation 2.2.1 (Presuppositions). With a more complex notion of context, it is more important than ever for us to decide over what Γ the judgment $\Gamma \vdash A \text{ type}$ ranges. We will say that the judgment $\Gamma \vdash A \text{ type}$ is only well-formed when $\vdash \Gamma \text{ cx}$ holds, as a matter of “meta-type discipline,” and similarly that the judgment $\Gamma \vdash a : A$ is only well-formed when $\Gamma \vdash A \text{ type}$ (and thus also $\vdash \Gamma \text{ cx}$).

One often says that $\vdash \Gamma \text{ cx}$ is a *presupposition* of the judgment $\Gamma \vdash A \text{ type}$, and that the judgments $\vdash \Gamma \text{ cx}$ and $\Gamma \vdash A \text{ type}$ are presuppositions of $\Gamma \vdash a : A$. We will globally adopt the convention that whenever we assert the truth of some judgment in prose or as the premise of a rule, we also implicitly assert that its presuppositions hold. Dually, we will be careful to check that none of our rules have meta-ill-typed conclusions.

Now that we have added a term variable context to the type well-formedness judgment, we can explain when $(x : A) \rightarrow B$ is a type: it is a (well-formed) type in Γ when A is a type in Γ and B is a type in $\Gamma, x : A$, as follows.

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma, x : A \vdash B \text{ type}}{\Gamma \vdash (x : A) \rightarrow B \text{ type}}$$

Rules like this describing how to create a type are known as *formation rules*, to parallel the terminology of introduction and elimination rules.

We can now sketch the formation rules for many of the types we encountered in Chapter 1. Dependent types like $_ \equiv _$ and \mathbf{Vec} are particularly interesting because they entangle the $\Gamma \vdash A$ type judgment with the term well-formedness judgment $\Gamma \vdash a : A$.

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{Nat} \text{ type}} \quad \frac{\Gamma \vdash A \text{ type} \quad \Gamma \vdash n : \mathbf{Nat}}{\Gamma \vdash \mathbf{Vec} A n \text{ type}} \quad \frac{\Gamma \vdash a : A \quad \Gamma \vdash b : A}{\Gamma \vdash a \equiv b \text{ type}}$$

Note that the convention of presuppositions outlined in Notation 2.2.1 means that the second and third rules have an implicit $\vdash \Gamma \text{ cx}$ premise, and the third rule also has an implicit $\Gamma \vdash A \text{ type}$ premise. To see that the conclusions of these rules are meta-well-typed, we must check that $\vdash \Gamma \text{ cx}$ holds in each case; this is an explicit premise of the first rule and a presupposition of the premises of the second and third rules.

The formation rule for propositional equality $_ \equiv _$ in particular is a major source of dependency because it singlehandedly allows arbitrary terms of arbitrary type to occur within types. In fact, this rule by itself causes the inference rules of all three judgments $\vdash \Gamma \text{ cx}$, $\Gamma \vdash A \text{ type}$, and $\Gamma \vdash a : A$ to all depend on one another pairwise.

Exercise 2.1. Attempt to derive that $(n : \mathbf{Nat}) \rightarrow \mathbf{Vec} \mathbf{String} (\text{succ } n)$ is a well-formed type in the empty context $\mathbf{1}$, using the rules introduced in this section thus far. Several rules are missing; which judgments can you not yet derive?

The variable rule Let us turn now to the term judgment $\Gamma \vdash a : A$, and in particular the rule stating that term variables in the context are well-formed terms. For simplicity, imagine the special case where the last variable is the one under consideration:

$$\frac{}{\Gamma, x : A \vdash x : A} \text{ !?}$$

This rule needs considerable work, as neither of the conclusion's presuppositions, $\vdash (\Gamma, x : A) \text{ cx}$ and $\Gamma, x : A \vdash A \text{ type}$, currently hold. We can address the former by adding premises $\vdash \Gamma \text{ cx}$ and $\Gamma \vdash A \text{ type}$ to the rule, from which it follows that $\vdash (\Gamma, x : A) \text{ cx}$.² As for the latter, note that $\Gamma \vdash A \text{ type}$ does not actually imply $\Gamma, x : A \vdash A \text{ type}$ —this would require proving a *weakening lemma* (see Lemma 2.1.6) for types! (Conversely, if the rule has the premise $\Gamma \vdash A \text{ type}$, then we cannot establish well-formedness of the context.)

There are several ways to proceed. One is to prove a weakening lemma, but given that the well-formedness of the variable rule requires weakening, it is necessary to prove all our well-formedness, weakening, and substitution lemmas by a rather heavy simultaneous induction. A second approach would be to add a silent weakening *rule* stating that

²Of course one could just directly add the premise $\vdash (\Gamma, x : A) \text{ cx}$, but our short-term memory is robust enough to recall that our next task is to ensure that A is a type.

$\Gamma, x : A \vdash B$ type whenever $\Gamma \vdash B$ type; however, this introduces ambiguity into our rules regarding the context(s) in which a type or term is well-formed.

We opt for a third option, which is to add *explicit* weakening rules asserting the existence of an operation sending types and terms in context Γ to types and terms in context $\Gamma, x : A$, both written $-[\mathbf{p}]$. (This notation will become less mysterious later.)

$$\frac{\Gamma \vdash B \text{ type} \quad \Gamma \vdash A \text{ type}}{\Gamma, x : A \vdash B[\mathbf{p}] \text{ type}} \qquad \frac{\Gamma \vdash b : B \quad \Gamma \vdash A \text{ type}}{\Gamma, x : A \vdash b[\mathbf{p}] : B[\mathbf{p}]}$$

Note that the type weakening rule is needed to make sense of the term weakening rule.

We can now fix the variable rule we wrote above: using $-[\mathbf{p}]$ to weaken A by itself, we move A from context Γ to $\Gamma, x : A$ as required in the conclusion of the rule.

$$\frac{\vdash \Gamma \text{ cx} \quad \Gamma \vdash A \text{ type}}{\Gamma, x : A \vdash x : A[\mathbf{p}]}$$

To use variables that occur earlier in the context, we can apply weakening repeatedly until they are the last variable. Suppose that $1 \vdash A$ type and $x : A \vdash B$ type, and in the context $x : A, y : B$ we want to use the variable x . Ignoring the $y : B$ in the context for a moment, we know that $x : A \vdash x : A[\mathbf{p}]$ by the last variable rule; thus by weakening we have $x : A, y : B \vdash x[\mathbf{p}] : A[\mathbf{p}][\mathbf{p}]$. In general, we can derive the following principle:

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma, x : A \vdash B_1 \text{ type} \quad \dots \quad \Gamma, x : A, y_1 : B_1, \dots \vdash B_n \text{ type}}{\Gamma, x : A, y_1 : B_1, \dots, y_n : B_n \vdash \underbrace{x[\mathbf{p}] \dots [\mathbf{p}]}_{n \text{ times}} : \underbrace{A[\mathbf{p}] \dots [\mathbf{p}]}_{n+1 \text{ times}}}$$

This approach to variables is elegant in that it breaks the standard variable rule into two simpler primitives: a rule for the last variable, and rules for type and term weakening. However, it introduces a redundancy in our notation, because the term $x[\mathbf{p}]^n$ encodes in two different ways the variable to which it refers: by the name x as well as positionally by the number of weakenings n .

A happy accident of our presentation of the variable rule is thus that we can delete variable names altogether; in Section 2.3 we will present contexts simply as lists of types $A.B.C$ with no variable names, and adopt a single notation for “the last variable in the context,” an encoding of the lambda calculus known as *de Bruijn indexing* [Bru72]. Conceptual elegance notwithstanding, this notation is very unfriendly to the reader in larger examples³ so we will continue to use named variables outside of the rules themselves; translating between the two notations is purely mechanical.

³According to Conor McBride, “Bob Atkey once memorably described the capacity to put up with de Bruijn indices as a Cylon detector.” (<https://mazzo.li/epilogue/index.html%3Fp=773.html>)

Remark 2.2.2. The first author wishes to mention another approach to maintaining readability, which is to continue using both named variables and explicit weakenings [Gra09]; this approach has the downside of requiring us to explain variable binding, but is simultaneously readable and precise about weakenings. \diamond

2.3 The calculus of substitutions

Weakening is one of two main operations in type theory that moves types and terms between contexts, the other being substitution of terms for variables. For the same reasons that we want to present weakening as an explicit type- and term-forming operation, we will also formulate substitution as an explicit operation subject to equations explicating how it computes on each construct of the theory.

However, rather than axiomatizing *single* substitutions and weakenings, we will axiomatize arbitrary compositions of substitutions and weakenings. In light of the fact that substitution shortens the context of a type/term and weakening lengthens it, these composite operations—called *simultaneous substitutions* (henceforth just substitutions)—can turn any context Γ into any other context Δ .

We thus add one final judgment to our presentation of type theory, $\Delta \vdash \gamma : \Gamma$ (“ γ is a substitution from Δ to Γ ”), corresponding to operations that send types/terms from context Γ to context Δ . (Not a typo; we will address the “backwards” notation later.)

Notation 2.3.1. Type theory has four basic judgments and three equality judgments:

1. $\vdash \Gamma \text{ cx}$ asserts that Γ is a context.
2. $\Delta \vdash \gamma : \Gamma$, presupposing $\vdash \Delta \text{ cx}$ and $\vdash \Gamma \text{ cx}$, asserts that γ is a substitution from Δ to Γ .
3. $\Gamma \vdash A \text{ type}$, presupposing $\vdash \Gamma \text{ cx}$, asserts that A is a type in context Γ .
4. $\Gamma \vdash a : A$, presupposing $\vdash \Gamma \text{ cx}$ and $\Gamma \vdash A \text{ type}$, asserts that a is an element/term of type A in context Γ .
- 2'. $\Delta \vdash \gamma = \gamma' : \Gamma$, presupposing $\Delta \vdash \gamma : \Gamma$ and $\Delta \vdash \gamma' : \Gamma$, asserts that γ, γ' are equal substitutions from Δ to Γ .
- 3'. $\Gamma \vdash A = A' \text{ type}$, presupposing $\Gamma \vdash A \text{ type}$ and $\Gamma \vdash A' \text{ type}$, asserts that A, A' are equal types in context Γ .
- 4'. $\Gamma \vdash a = a' : A$, presupposing $\Gamma \vdash a : A$ and $\Gamma \vdash a' : A$, asserts that a, a' are equal elements of type A in context Γ .

Notation 2.3.2. We write Cx for the set of contexts, $Sb(\Delta, \Gamma)$ for the set of substitutions from Δ to Γ , $Ty(\Gamma)$ for the set of types in context Γ , and $Tm(\Gamma, A)$ for the set of terms of type A in context Γ .

This presentation of dependent type theory is known as the *substitution calculus* [ML92; Tas93]. Perhaps unsurprisingly, we must discuss a considerable number of rules governing substitutions before presenting any concrete type and term formers; we devote this section to those rules, and cover the main connectives of type theory in Section 2.4.

Contexts The rules for contexts are as in Section 2.2, but without variable names:

$$\frac{}{\vdash 1 \text{ cx}} \qquad \frac{\vdash \Gamma \text{ cx} \quad \Gamma \vdash A \text{ type}}{\vdash \Gamma.A \text{ cx}}$$

Although there is no context equality judgment, note that two contexts *can* be equal without being syntactically identical. If $1 \vdash A = A' \text{ type}$ then $1.A$ and $1.A'$ are equal contexts on the basis that, like all operations of the theory, context extension respects equality in both arguments. We have omitted the $\vdash \Gamma = \Gamma' \text{ cx}$ judgment for the simple reason that there would be no rules governing it: the only reason why two contexts can be equal is that their types are pairwise equal.

Substitutions The purpose of a substitution $\Delta \vdash \gamma : \Gamma$ is to shift types and terms from context Γ to context Δ :

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type}}{\Delta \vdash A[\gamma] \text{ type}} \qquad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : A}{\Delta \vdash a[\gamma] : A[\gamma]}$$

Unlike the substitution operation of Section 2.1, which was a function on terms defined by cases, these rules define two binary type- and term- forming operations that take a type (resp., term) and a substitution as input and produce a new type (resp., term). Note also that, despite sharing a notation, type and term substitution are two distinct operations.

The simplest interesting substitution is weakening, written \mathbf{p} :⁴

$$\frac{\Gamma \vdash A \text{ type}}{\Gamma.A \vdash \mathbf{p} : \Gamma}$$

In concert with the substitution rules above we can recover the weakening rules from the previous section, e.g., if $\Gamma \vdash B \text{ type}$ and $\Gamma \vdash A \text{ type}$ then $\Gamma, x : A \vdash B[\mathbf{p}] \text{ type}$.

⁴This mysterious name can be explained by the fact that weakening corresponds semantically to a projection map; \mathbf{p} can thus be pronounced as either “weakening” or “projection”.

Because substitutions $\Delta \vdash \gamma : \Gamma$ encode arbitrary compositions of context-shifting operations, we also have rules that close substitutions under nullary and binary composition:

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{id} : \Gamma} \qquad \frac{\Gamma_2 \vdash \gamma_1 : \Gamma_1 \quad \Gamma_1 \vdash \gamma_0 : \Gamma_0}{\Gamma_2 \vdash \gamma_0 \circ \gamma_1 : \Gamma_0}$$

These operations are unital and associative as one might expect:

$$\frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \gamma \circ \mathbf{id} = \mathbf{id} \circ \gamma = \gamma : \Gamma} \qquad \frac{\Gamma_3 \vdash \gamma_2 : \Gamma_2 \quad \Gamma_2 \vdash \gamma_1 : \Gamma_1 \quad \Gamma_1 \vdash \gamma_0 : \Gamma_0}{\Gamma_3 \vdash \gamma_0 \circ (\gamma_1 \circ \gamma_2) = (\gamma_0 \circ \gamma_1) \circ \gamma_2 : \Gamma_0}$$

We can summarize the rules above by stating that there is a *category* whose objects are contexts and whose morphisms are substitutions.

We have already seen that substitutions shift the contexts of types and terms by $-[\gamma]$; they also shift the context of other substitutions by precomposition. Later we will have occasion to discuss all three context-shifting functions between sorts that are induced by substitutions, as follows.

Notation 2.3.3. Given a substitution $\Delta \vdash \gamma : \Gamma$, we write γ^* for the following functions:

- $\xi \mapsto \xi \circ \gamma : \text{Sb}(\Gamma, \Xi) \rightarrow \text{Sb}(\Delta, \Xi)$,
- $A \mapsto A[\gamma] : \text{Ty}(\Gamma) \rightarrow \text{Ty}(\Delta)$, and
- $a \mapsto a[\gamma] : \text{Tm}(\Gamma, A) \rightarrow \text{Tm}(\Delta, A[\gamma])$.

Composite substitutions introduce a possible redundancy into our rules: what is the difference between substituting by γ_0 and then by γ_1 versus substituting once by $\gamma_0 \circ \gamma_1$? We add equations asserting that substituting by \mathbf{id} is the identity and substituting by a composite is composition of substitutions:

$$\frac{\Gamma \vdash A \text{ type}}{\Gamma \vdash A[\mathbf{id}] = A \text{ type}} \qquad \frac{\Gamma \vdash a : A}{\Gamma \vdash a[\mathbf{id}] = a : A}$$

$$\frac{\Gamma_2 \vdash \gamma_1 : \Gamma_1 \quad \Gamma_1 \vdash \gamma_0 : \Gamma_0 \quad \Gamma_0 \vdash A \text{ type}}{\Gamma_2 \vdash A[\gamma_0 \circ \gamma_1] = A[\gamma_0][\gamma_1] \text{ type}} \qquad \frac{\Gamma_2 \vdash \gamma_1 : \Gamma_1 \quad \Gamma_1 \vdash \gamma_0 : \Gamma_0 \quad \Gamma_0 \vdash a : A}{\Gamma_2 \vdash a[\gamma_0 \circ \gamma_1] = a[\gamma_0][\gamma_1] : A[\gamma_0 \circ \gamma_1]}$$

We can summarize the rules above by stating that the γ^* operations respect identity and composition of substitutions, or more compactly, that the collections of types and terms form *presheaves* $\text{Ty}(-)$ and $\sum_{A:\text{Ty}(-)} \text{Tm}(-, A)$ on the category of contexts, with restriction maps given by substitution (a perspective which inspires the notation γ^*).

Before moving on, it is instructive to once again convince ourselves that the rules above are meta-well-typed. In particular, the conclusion of the second rule is only sensible if $\Gamma \vdash a[\mathbf{id}] : A$, but according to the rule for term substitution we only have $\Gamma \vdash a[\mathbf{id}] : A[\mathbf{id}]$. To make sense of this rule we must refer to the previous rule equating the types $A[\mathbf{id}]$ and A . A consequence of this type equation is that terms of type $A[\mathbf{id}]$ are equivalently terms of type A ,⁵ and thus $\Gamma \vdash a[\mathbf{id}] : A$ as required. This is a paradigmatic example of the deeply intertwined nature of the rules of dependent type theory; in particular, *we cannot defer equations* to the end of our construction the way we did in Section 2.1 because many rules are only sensible after imposing certain equations.

The variable rule revisited As in the previous section, the variable rule is restricted to the last entry in the context, which we (unambiguously) always name \mathbf{q} .⁶

$$\frac{\Gamma \vdash A \text{ type}}{\Gamma.A \vdash \mathbf{q} : A[\mathbf{p}]}$$

Writing \mathbf{p}^n for the n -fold composition of \mathbf{p} with itself (with $\mathbf{p}^0 = \mathbf{id}$), the following rule is *derivable* from other rules (notated \Rightarrow) and thus not explicitly included in our system:

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B_1 \text{ type} \quad \dots \quad \Gamma.A.B_1 \dots \vdash B_n \text{ type}}{\Gamma.A.B_1 \dots B_n \vdash \mathbf{q}[\mathbf{p}^n] : A[\mathbf{p}^{n+1}]} \Rightarrow$$

Thus a variable in our system is a term of the form $\mathbf{q}[\mathbf{p}^n]$, where n is its de Bruijn index.

Terminal substitutions Our notation $\Delta \vdash \gamma : \Gamma$ for substitutions is no accident; it is indeed a good mental model to think of such substitutions as “terms of type Γ in context Δ .” To understand why, let us think back to propositional logic. A term $\mathbf{1}.B \vdash c : C$ can be seen as a proof of C under the hypothesis B , i.e., a proof that $B \implies C$. Given a substitution $\mathbf{1}.A \vdash b : \mathbf{1}.B$ we can obtain a term $\mathbf{1}.A \vdash c[b] : C[b]$, or a proof that $A \implies C$. This suggests that substituting corresponds logically to a “cut,” and b to a proof that $A \implies B$.

Returning to the general case, contexts are lists of hypotheses, and a substitution $\Delta \vdash \gamma : \Gamma$ states that we can prove all the hypotheses of Γ using the hypotheses of Δ . Thus anything that is true under the hypotheses Γ is also true under the hypotheses Δ —hence the contravariance of the substitution operation.

More concretely, the idea is that a substitution $\Delta \vdash \gamma : \mathbf{1}.A_1 \dots A_n$ is an n -tuple of terms a_1, \dots, a_n of types A_1, \dots, A_n , all in context Δ , and applying the substitution γ has

⁵In some presentations of type theory this principle is explicit and is known as the *type conversion rule*. For us it is a consequence of the judgments respecting equality, i.e., $\text{Tm}(\Gamma, A[\mathbf{id}]) = \text{Tm}(\Gamma, A)$ as sets.

⁶This mysterious name is chosen to pair well with the name \mathbf{p} that we gave weakening; \mathbf{q} can thus be pronounced as either “variable” or “qvariable”.

the effect of substituting a_1 for the first variable, a_2 for the second variable, ... and a_n for the last variable. The final subtlety is that each type A_i is in general dependent on all the previous A_j for $j < i$, so the type of a_2 is not just A_2 but " $A_2[a_1/x_1]$," so to speak, all the way through " $a_n : A_n[a_1/x_1, \dots, a_{n-1}/x_{n-1}]$."

If all of this sounds very complicated, well... at any rate, the remaining rules governing substitution define such n -tuples in two cases, 0 and $n + 1$. The nullary case is fairly simple: any substitution $\Gamma \vdash \delta : \mathbf{1}$ into the empty context (a length-zero list of types) is necessarily the empty tuple $\langle \rangle$, which we spell !.

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash ! : \mathbf{1}} \qquad \frac{\Gamma \vdash \delta : \mathbf{1}}{\Gamma \vdash ! = \delta : \mathbf{1}}$$

These rules state that $\mathbf{1}$ is a terminal object in the category of contexts, a perspective which inspires the notations $\mathbf{1}$ and !.

Substitution extension The other case concerns substitutions $\Delta \vdash - : \Gamma.A$ into a context extension. Recall that $\Gamma.A$ is an $(n + 1)$ -tuple of types when Γ is an n -tuple of types, and suppose that $\Delta \vdash \gamma : \Gamma$, which is to say that γ is an n -tuple of terms (in context Δ) whose types are those in Γ . To extend this n -tuple to an $(n + 1)$ -tuple of terms whose types are those in $\Gamma.A$, we simply adjoin one more term a in context Δ with type $A[\gamma]$, where this substitution plugs the n previously-given terms into the dependencies of A .

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type} \quad \Delta \vdash a : A[\gamma]}{\Delta \vdash \gamma.a : \Gamma.A}$$

The final three rules of our calculus are equations governing this substitution former:

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type} \quad \Delta \vdash a : A[\gamma]}{\Delta \vdash \mathbf{p} \circ (\gamma.a) = \gamma : \Gamma} \qquad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type} \quad \Delta \vdash a : A[\gamma]}{\Delta \vdash \mathbf{q}[\gamma.a] = a : A[\gamma]}$$

$$\frac{\Gamma \vdash A \text{ type} \quad \Delta \vdash \gamma : \Gamma.A}{\Delta \vdash \gamma = (\mathbf{p} \circ \gamma).\mathbf{q}[\gamma] : \Gamma.A}$$

Imagining for the moment that $\Gamma = x_1 : A_1, \dots, x_n : A_n$ and $\gamma = [a_1/x_1, \dots, a_n/x_n]$, the second rule states that $x_n[a_1/x_1, \dots, a_n/x_n] = a_n$, in other words, that substituting into the last variable x_n replaces that variable by the last term a_n . The first rule states in essence that substituting into a type/term that does not mention (is weakened by) x_n is the same as dropping the last term a_n/x_n from the substitution, i.e., $[a_1/x_1, \dots, a_{n-1}/x_{n-1}]$.

Finally, the third rule states that every substitution γ into the context $\Gamma.A$ is of the form $\gamma_0.a$, where a is determined by the behavior of γ on the last variable, and γ_0 is determined by the behavior of γ on the first n variables. (See Exercise 2.5.)

All of these rules in this section determine a category (of contexts and substitutions) with extra structure, known collectively as a *category with families* [Dyb96]. We will refer to any system that extends this collection of rules as a *Martin-Löf type theory*.

Exercise 2.2. Show that substitutions $\Gamma \vdash \gamma : \Gamma.A$ satisfying $\mathbf{p} \circ \gamma = \mathbf{id}$ are in bijection with terms $\Gamma \vdash a : A$.

Exercise 2.3. Show that $(\gamma.a) \circ \delta = (\gamma \circ \delta).a[\delta]$.

Exercise 2.4. Given $\Delta \vdash \gamma : \Gamma$ and $\Gamma \vdash A$ type, construct a substitution that we will name $\gamma.A$, satisfying $\Delta.A[\gamma] \vdash \gamma.A : \Gamma.A$.

Exercise 2.5. Suppose that $\Gamma \vdash A$ type and $\vdash \Delta$ cx. Show that substitutions $\Delta \vdash \gamma : \Gamma.A$ are in bijection with pairs of a substitution $\Delta \vdash \gamma_0 : \Gamma$ and a term $\Delta \vdash a : A[\gamma_0]$.

2.4 Internalizing judgmental structure: $\Pi, \Sigma, \text{Eq}, \text{Unit}$

With the basic structure of dependent type theory finally out of the way, we are prepared to define standard type and term formers, starting with the best-behaved connectives: dependent products, dependent sums, extensional equality, and the unit type. Unlike inductive types (Section 2.5), each of these connectives can be described concisely as internalizing judgmental structure of some kind.

2.4.1 Dependent products

We start with dependent function types, also known as *dependent products* or Π -types. The formation rule is as in Section 2.2, but without variable names:⁷

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type}}{\Gamma \vdash \Pi(A, B) \text{ type}}$$

Remark 2.4.1. The Π notation and terminology is inspired by this type corresponding semantically to a set-indexed product of sets $\prod_{a \in A} B_a$. Indexed products generalize ordinary products in the sense that $\prod_{a \in \{1,2\}} B_a \cong B_1 \times B_2$. \diamond

Remarkably, the substitution calculus ensures that these rules are almost indistinguishable from the introduction and elimination rules of simple function types in Section 2.1,

⁷We have switched our notation from $(x : A) \rightarrow B$ because it is awkward without named variables.

with some minor additional bookkeeping to move types to the appropriate contexts:

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash b : B}{\Gamma \vdash \lambda(b) : \Pi(A, B)} \quad \frac{\Gamma \vdash a : A \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash f : \Pi(A, B)}{\Gamma \vdash \mathbf{app}(f, a) : B[\mathbf{id}.a]}$$

There continue to be a few notational shifts: λ s no longer come with variable names, and we write $\mathbf{app}(f, a)$ rather than $f a$ just to emphasize that function application is a term constructor. The reader should convince themselves that in the final rule, $\Gamma \vdash B[\mathbf{id}.a]$ type; this substitutes a for the last variable in B , leaving the rest of the context unchanged.

Next we must specify equations not only on the introduction and elimination forms, but on the type former itself. There are two groups of equations we must impose; the first group explains how substitutions act on all three of these operations:

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type}}{\Delta \vdash \Pi(A, B)[\gamma] = \Pi(A[\gamma], B[\gamma.A]) \text{ type}} \quad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type} \quad \Gamma.A \vdash b : B}{\Delta \vdash \lambda(b)[\gamma] = \lambda(b[\gamma.A]) : \Pi(A, B)[\gamma]}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : A \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash f : \Pi(A, B)}{\Delta \vdash \mathbf{app}(f, a)[\gamma] = \mathbf{app}(f[\gamma], a[\gamma]) : B[\gamma.a[\gamma]]}$$

Roughly speaking, these three rules state that substitutions commute past each type and term former, but B and b are well-formed in a larger context ($\Gamma.A$) than the surrounding term (Γ), requiring us to “shift” the substitution so that it leaves the bound variable of type A unchanged while continuing to act on all the free variables in Γ . (The “shifted” substitution $\gamma.A$ in these rules is the derived form defined in Exercise 2.4.)

Once again we should pause and convince ourselves that these rules are meta-well-typed. Echoing the phenomenon we saw in Section 2.3 with $\Gamma \vdash a[\mathbf{id}] : A$, we need to use the substitution rule for $\Pi(A, B)[\gamma]$ to see that the right-hand side of the substitution rules for $\lambda(b)[\gamma]$ and $\mathbf{app}(f, a)[\gamma]$ are well-typed.

Exercise 2.6. Check that the substitution rule for $\mathbf{app}(f, a)[\gamma]$ is meta-well-typed; in particular, show that both $\mathbf{app}(f, a)[\gamma]$ and $\mathbf{app}(f[\gamma], a[\gamma])$ have the type $B[\gamma.a[\gamma]]$.

This pattern will continue: every time we introduce a new type or term former θ , we will add an equation $\theta(a_1, \dots, a_n)[\gamma] = \theta(a_1[\gamma_1], \dots, a_n[\gamma_n])$ stating that substitutions push past θ , adjusted as necessary in each argument. These rules are quite mechanical and can even be automatically derived in some frameworks, but they are at the heart of type theory itself. From a logical perspective, they ensure that quantifier instantiation is uniform. From a mathematical perspective, as we will see in Section 2.4.2, they assert the naturality of type-theoretic constructions. And from an implementation perspective, these rules can be assembled into a substitution algorithm, ensuring that substitutions can be computed automatically by proof assistants.

Remark 2.4.2. The difference between this approach to substitution and the one outlined in Section 2.1 is one of *derivability* vs *admissibility*. In the simply-typed setting, the fact that all terms enjoy substitution is not part of the system but rather must be proven (and even constructed in the first place) by induction over the structure of terms, and so adding new constructs to the theory may cause substitution to fail.

In the substitution calculus, we assert that all types and terms enjoy substitution as basic rules of the theory, and later add equations specifying how substitution computes; thus any extension of the theory is guaranteed to enjoy substitution. Because substitution is a crucial aspect of dependent type theory, we find this latter approach more ergonomic. \diamond

The second group of equations is the β - and η -rules introduced in Section 2.1, completing our presentation of dependent product types.

$$\frac{\Gamma \vdash a : A \quad \Gamma.A \vdash b : B}{\Gamma \vdash \mathbf{app}(\lambda(b), a) = b[\mathbf{id}.a] : B[\mathbf{id}.a]} \quad \frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash f : \Pi(A, B)}{\Gamma \vdash f = \lambda(\mathbf{app}(f[\mathbf{p}], \mathbf{q})) : \Pi(A, B)}$$

Exercise 2.7. Carefully explain why the η -rule above is meta-well-typed, in particular why $\lambda(\mathbf{app}(f[\mathbf{p}], \mathbf{q}))$ has the right type. Explicitly point out all the other rules and equations (e.g., Π -introduction, Π -elimination, weakening) to which you refer.

Exercise 2.8. Show that using Π -types we can define a non-dependent function type whose formation rule states that if $\Gamma \vdash A \text{ type}$ and $\Gamma \vdash B \text{ type}$ then $\Gamma \vdash A \rightarrow B \text{ type}$. Then define the introduction and elimination rules from Section 2.1 for this encoding, and check that the β - and η -rules from Section 2.1 hold. (Hint: it is incorrect to define $A \rightarrow B := \Pi(A, B)$.)

Exercise 2.9. As discussed in Section 2.3, two contexts that are not syntactically identical may nevertheless be equal. Give an example.

2.4.2 Dependent products internalize hypothetical judgments

With one type constructor, two term constructors, and five equations, it is natural to wonder whether we have written “enough” or “the correct” rules to specify Π -types. One may also wonder whether there is an easier way. We now introduce a methodology for making sense of this collection of rules, and show how we can use this methodology to more efficiently define the later connectives. In short, we will view connectives as *internalizations of judgmental structure*, and $\Gamma \vdash - : \Pi(A, B)$ in particular as an internalization of the hypothetical judgment $\Gamma.A \vdash - : B$.

Remark 2.4.3. In this book we limit ourselves to a semi-informal discussion of this perspective, which can be made fully precise with the language of category theory. For

instance, using the framework of natural models, Awodey [Awo18] shows that the rules above exactly capture that Π -types classify the hypothetical judgment in a precise sense. \diamond

Analyzing context extension To warm up, let us begin by recalling Exercise 2.5, which establishes the following bijection of sets for every Δ, Γ , and A :

$$\{\gamma \mid \Delta \vdash \gamma : \Gamma.A\} \cong \{(\gamma_0, a) \mid \Delta \vdash \gamma_0 : \Gamma \wedge \Delta \vdash a : A[\gamma_0]\}$$

Using Notation 2.3.2 we equivalently write:

$$\iota_{\Delta, \Gamma, A} : \text{Sb}(\Delta, \Gamma.A) \cong \sum_{\gamma \in \text{Sb}(\Delta, \Gamma)} \text{Tm}(\Delta, A[\gamma])$$

where $\sum_{a \in A} B_a$ is our notation for the set-indexed coproduct of sets $\coprod_{a \in A} B_a$.

As stated, the bijections $\iota_{\Delta, \Gamma, A}$ and $\iota_{\Delta', \Gamma', A'}$ may be totally unrelated, but it turns out that this collection of bijections is actually *natural* (or “parametric”) in Δ in the sense that the behavior of $\iota_{\Delta_0, \Gamma, A}$ and $\iota_{\Delta_1, \Gamma, A}$ are correlated when we have a substitution from Δ_0 to Δ_1 .

Because these bijections have different types, to make this idea precise we must find a way to relate their differing domains $\text{Sb}(\Delta_0, \Gamma.A)$ and $\text{Sb}(\Delta_1, \Gamma.A)$ with one another, as well as their codomains $\sum_{\gamma \in \text{Sb}(\Delta_0, \Gamma)} \text{Tm}(\Delta_0, A[\gamma])$ and $\sum_{\gamma \in \text{Sb}(\Delta_1, \Gamma)} \text{Tm}(\Delta_1, A[\gamma])$.

We have already seen the former in Notation 2.3.3: every substitution $\Delta_0 \vdash \delta : \Delta_1$ induces a function $\delta^* : \text{Sb}(\Delta_1, \Gamma.A) \rightarrow \text{Sb}(\Delta_0, \Gamma.A)$. We leave the latter as an exercise:

Exercise 2.10. Given $\Delta_0 \vdash \delta : \Delta_1$, use δ^* (Notation 2.3.3) to define the following function:

$$\sum_{\delta^*} \delta^* : \sum_{\gamma \in \text{Sb}(\Delta_1, \Gamma)} \text{Tm}(\Delta_1, A[\gamma]) \rightarrow \sum_{\gamma \in \text{Sb}(\Delta_0, \Gamma)} \text{Tm}(\Delta_0, A[\gamma])$$

Proof. Define $(\sum_{\delta^*} \delta^*)(\gamma, a) = (\delta^* \gamma, \delta^* a) = (\gamma \circ \delta, a[\delta])$. \square

With these functions in hand we can now explain precisely what we mean by the naturality of $\iota_{-, \Gamma, A}$. Fix a substitution $\Delta_0 \vdash \delta : \Delta_1$. We have two different ways of turning a substitution $\Delta_1 \vdash \gamma : \Gamma.A$ into an element of $\sum_{\gamma_0 \in \text{Sb}(\Delta_0, \Gamma)} \text{Tm}(\Delta_0, A[\gamma_0])$, depicted by the “right then down” and “down then right” paths in the diagram below:

$$\begin{array}{ccc} \text{Sb}(\Delta_1, \Gamma.A) & \xrightarrow{\iota_{\Delta_1, \Gamma, A}} & \sum_{\gamma \in \text{Sb}(\Delta_1, \Gamma)} \text{Tm}(\Delta_1, A[\gamma]) \\ \delta^* \downarrow & & \downarrow \sum_{\delta^*} \delta^* \\ \text{Sb}(\Delta_0, \Gamma.A) & \xrightarrow{\iota_{\Delta_0, \Gamma, A}} & \sum_{\gamma \in \text{Sb}(\Delta_0, \Gamma)} \text{Tm}(\Delta_0, A[\gamma]) \end{array}$$

Going “right then down” we obtain

$$\begin{array}{ccc} \gamma & \longmapsto & \iota_{\Delta_1, \Gamma, A}(\gamma) \\ & & \downarrow \\ & & (\sum_{\delta^*} \delta^*)(\iota_{\Delta_1, \Gamma, A}(\gamma)) \end{array}$$

and going “down then right” we obtain $\gamma \mapsto \gamma \circ \delta \mapsto \iota_{\Delta_0, \Gamma, A}(\gamma \circ \delta)$.

We say that the family of isomorphisms $\Delta \mapsto \iota_{\Delta, \Gamma, A}$ is natural when these two paths always yield the same result, i.e., when $(\sum_{\delta^*} \delta^*)(\iota_{\Delta_1, \Gamma, A}(\gamma)) = \iota_{\Delta_0, \Gamma, A}(\gamma \circ \delta)$ for every $\Delta_0 \vdash \delta : \Delta_1$ and γ . In other words, $\iota_{\Delta_0, \Gamma, A}$ and $\iota_{\Delta_1, \Gamma, A}$ “do the same thing” as soon as you correct the mismatch in their types by pre- and post-composing the appropriate maps.

Exercise 2.11. Prove that ι is natural, i.e., that the following maps are equal:

$$\sum_{\delta^*} \delta^* \circ \iota_{\Delta_1, \Gamma, A} = \iota_{\Delta_0, \Gamma, A} \circ \delta^* : \text{Sb}(\Delta_1, \Gamma.A) \rightarrow \sum_{\gamma \in \text{Sb}(\Delta_0, \Gamma)} \text{Tm}(\Delta_0, A[\gamma])$$

Proof. Suppose $\gamma \in \text{Sb}(\Delta_1, \Gamma.A)$. Unfolding the solutions to Exercises 2.5 and 2.10,

$$\begin{aligned} (\sum_{\delta^*} \delta^*)(\iota_{\Delta_1, \Gamma, A}(\gamma)) &= (\sum_{\delta^*} \delta^*)(\mathbf{p} \circ \gamma, \mathbf{q}[\gamma]) = ((\mathbf{p} \circ \gamma) \circ \delta, \mathbf{q}[\gamma][\delta]) \\ \iota_{\Delta_0, \Gamma, A}(\delta^*(\gamma)) &= \iota_{\Delta_0, \Gamma, A}(\gamma \circ \delta) = (\mathbf{p} \circ (\gamma \circ \delta), \mathbf{q}[\gamma \circ \delta]) \end{aligned}$$

which are equal by the functoriality of substitution. \square

The terminology of “natural” comes from category theory, where $\iota_{-, \Gamma, A}$ is known as a natural isomorphism, but we will prove and use naturality conditions without referring to the general concept. One useful consequence of naturality is the following:

Exercise 2.12. *Without* unfolding the definition of ι , show that the naturality of ι and the fact that $\iota_{\Delta, \Gamma, A}$ and $\iota_{\Delta, \Gamma, A}^{-1}$ are inverses together imply that ι^{-1} is natural, i.e., that

$$\iota_{\Delta_0, \Gamma, A}^{-1} \circ \sum_{\delta^*} \delta^* = \delta^* \circ \iota_{\Delta_1, \Gamma, A}^{-1} : \sum_{\gamma \in \text{Sb}(\Delta_1, \Gamma)} \text{Tm}(\Delta_1, A[\gamma]) \rightarrow \text{Sb}(\Delta_0, \Gamma.A)$$

Proof. Apply $\iota_{\Delta_0, \Gamma, A}^{-1} \circ - \circ \iota_{\Delta_1, \Gamma, A}^{-1}$ to both sides of the naturality equation for ι and cancel:

$$\begin{aligned} \iota_{\Delta_0, \Gamma, A}^{-1} \circ \sum_{\delta^*} \delta^* \circ \iota_{\Delta_1, \Gamma, A} \circ \iota_{\Delta_1, \Gamma, A}^{-1} &= \iota_{\Delta_0, \Gamma, A}^{-1} \circ \iota_{\Delta_0, \Gamma, A} \circ \delta^* \circ \iota_{\Delta_1, \Gamma, A}^{-1} \\ \iota_{\Delta_0, \Gamma, A}^{-1} \circ \sum_{\delta^*} \delta^* &= \delta^* \circ \iota_{\Delta_1, \Gamma, A}^{-1} \end{aligned} \quad \square$$

Exercise 2.13. For categorically-minded readers: argue that ι is a natural isomorphism in the standard sense, by rephrasing Exercises 2.10 and 2.11 in terms of categories and functors.

Rather than defining context extension by the collection of rules in Section 2.3 and then characterizing it in terms of ι after the fact, we can actually define it directly as “a context $\Gamma.A$ for which $\text{Sb}(-, \Gamma.A)$ is naturally isomorphic to $\sum_{\gamma \in \text{Sb}(-, \Gamma)} \text{Tm}(-, A[\gamma])$,” which unfolds to all of the relevant rules.

In addition to its brevity, the true advantage of such characterizations is that they are less likely to “miss” some important aspect of the definition. Zooming out, this definition states that substitutions into $\Gamma.A$ are dependent pairs of a substitution γ into Γ and a term in $A[\gamma]$, which is exactly the informal description we started with in Section 2.3.

With that in mind, our program for justifying the rules of type theory is as follows:

Slogan 2.4.4. *A connective in type theory is given by (1) a natural type-forming operation and (2) a natural isomorphism relating that type’s terms to judgmentally-determined structure.*

We must unfortunately remain vague here about the meaning of “judgmentally-determined structure,” but it refers to sets constructed from the sorts $\text{Sb}(\Delta, \Gamma)$, $\text{Ty}(\Gamma)$, and $\text{Tm}(\Gamma, A)$ using natural operations such as dependent products and dependent sums—operations that are implicit in the meaning of inference rules. To make this more precise requires a formal treatment of the algebra of judgments via *logical frameworks*.

In addition, although this slogan will make quick work of the remainder of Section 2.4, we will need to revise it in Sections 2.5 and 2.6.

Π -types The rules in Section 2.4.1 precisely capture the existence of an operation

$$\Pi_{\Gamma} : (\sum_{A \in \text{Ty}(\Gamma)} \text{Ty}(\Gamma.A)) \rightarrow \text{Ty}(\Gamma)$$

natural in Γ (that is, one which commutes with substitution) along with the following family of isomorphisms also natural in Γ :

$$\iota_{\Gamma, A, B} : \text{Tm}(\Gamma, \Pi(A, B)) \cong \text{Tm}(\Gamma.A, B)$$

The first point expresses the formation rule and $\Pi(A, B)[\gamma] = \Pi(A[\gamma], B[\gamma.A])$. We focus on the second point, which characterizes the remaining rules in Section 2.4.1.

The reverse map $\iota_{\Gamma, A, B}^{-1} : \text{Tm}(\Gamma.A, B) \rightarrow \text{Tm}(\Gamma, \Pi(A, B))$ is the introduction rule, which sends terms $\Gamma.A \vdash b : B$ to $\lambda(b)$. The forward map is slightly more involved, but we can guess that it should correspond to elimination. In fact it is *application to a fresh variable*, or a combination of weakening and application—given $\Gamma \vdash f : \Pi(A, B)$, we weaken to $\Gamma.A \vdash f[\mathbf{p}] : \Pi(A, B)[\mathbf{p}]$ and then apply to \mathbf{q} , obtaining $\Gamma.A \vdash \mathbf{app}(f[\mathbf{p}], \mathbf{q}) : B$.

To complete this natural isomorphism we must check that it is an isomorphism, and that it is natural. We begin with the isomorphism: for all $\vdash \Gamma \text{ cx}$, $\Gamma \vdash A$ type, and $\Gamma.A \vdash B$ type,

$$\begin{aligned} \iota_{\Gamma,A,B}^{-1}(\iota_{\Gamma,A,B}(f)) &= f \\ \iota_{\Gamma,A,B}(\iota_{\Gamma,A,B}^{-1}(b)) &= b \end{aligned}$$

Unfolding definitions, we see that this isomorphism boils down essentially to β and η .

$$\begin{aligned} &\iota_{\Gamma,A,B}^{-1}(\iota_{\Gamma,A,B}(f)) \\ &= \lambda(\mathbf{app}(f[\mathbf{p}], \mathbf{q})) \\ &= f && \text{by the } \eta \text{ rule} \\ \\ &\iota_{\Gamma,A,B}(\iota_{\Gamma,A,B}^{-1}(b)) \\ &= \mathbf{app}(\lambda(b)[\mathbf{p}], \mathbf{q}) \\ &= \mathbf{app}(\lambda(b[\mathbf{p}.A]), \mathbf{q}) && \lambda(-) \text{ commutes with substitution} \\ &= b[\mathbf{p}.A \circ \mathbf{id}.\mathbf{q}] && \text{by the } \beta \text{ rule} \\ &= b[\mathbf{p}.\mathbf{q}] && \text{by Exercise 2.14 below} \\ &= b[\mathbf{id}] \\ &= b \end{aligned}$$

Exercise 2.14. Using the definition of $\mathbf{p}.A$ from Exercise 2.4, prove the substitution equality needed to complete the equational reasoning above.

As for the naturality of the isomorphisms ι , as before we must first explain how to relate the types of $\iota_{\Gamma,A,B}$ and $\iota_{\Delta,A[\gamma],B[\gamma.A]}$ given a substitution $\Delta \vdash \gamma : \Gamma$. In this case, the comparison functions are the following:

$$\begin{aligned} \gamma^* &: \text{Tm}(\Gamma, \Pi(A, B)) \rightarrow \text{Tm}(\Delta, \Pi(A[\gamma], B[\gamma.A])) \\ \gamma.A^* &: \text{Tm}(\Gamma.A, B) \rightarrow \text{Tm}(\Delta.A[\gamma], B[\gamma.A]) \end{aligned}$$

Naturality therefore states that “right then down” and “down then right” are equal in the following diagram. (By the reader’s argument in Exercise 2.12, naturality of ι automatically implies the naturality of ι^{-1} .)

$$\begin{array}{ccc} \text{Tm}(\Gamma, \Pi(A, B)) & \xrightarrow{\iota_{\Gamma,A,B}} & \text{Tm}(\Gamma.A, B) \\ \downarrow \gamma^* & & \downarrow \gamma.A^* \\ \text{Tm}(\Delta, \Pi(A[\gamma], B[\gamma.A])) & \xrightarrow{\iota_{\Delta,A[\gamma],B[\gamma.A]}} & \text{Tm}(\Delta.A[\gamma], B[\gamma.A]) \end{array}$$

Fixing $\Gamma \vdash f : \Pi(A, B)$, we show $\iota_{\Gamma, A, B}(f)[\gamma.A] = \iota_{\Delta, A[\gamma], B[\gamma.A]}(f[\gamma])$ by computing:

$$\begin{aligned}
& \iota_{\Gamma, A, B}(f)[\gamma.A] \\
&= \mathbf{app}(f[\mathbf{p}], \mathbf{q})[\gamma.A] \\
&= \mathbf{app}(f[\mathbf{p}][\gamma.A], \mathbf{q}[\gamma.A]) \quad \mathbf{app}(-, -) \text{ commutes with substitution} \\
&= \mathbf{app}(f[\mathbf{p} \circ \gamma.A], \mathbf{q}) \\
&= \mathbf{app}(f[\gamma \circ \mathbf{p}], \mathbf{q}) \\
& \iota_{\Delta, A[\gamma], B[\gamma.A]}(f[\gamma]) \\
&= \mathbf{app}(f[\gamma][\mathbf{p}], \mathbf{q}) \\
&= \mathbf{app}(f[\gamma \circ \mathbf{p}], \mathbf{q})
\end{aligned}$$

Thus all of the rules of Π -types can be summed up by a natural operation Π_{Γ} (formation and its substitution law) along with a natural isomorphism $\iota_{\Gamma, A, B} : \text{Tm}(\Gamma, \Pi(A, B)) \cong \text{Tm}(\Gamma.A, B)$ where ι^{-1} and ι are introduction and elimination, the round-trips are β and η , and naturality is the remaining substitution laws.

An alternative eliminator There is a strange asymmetry in the two maps ι and ι^{-1} underlying our natural isomorphism: the latter is literally the introduction rule, but the former combines elimination with weakening and the variable rule. It turns out that there is an equivalent formulation of Π -elimination more faithful to our current perspective:

$$\frac{\Gamma \vdash f : \Pi(A, B)}{\Gamma.A \vdash \lambda^{-1}(f) : B} \Rightarrow$$

Such a presentation replaces the current $\mathbf{app}(-, -)$, β , and η rules with the above rule along with new versions of β and η stating simply that $\lambda(\lambda^{-1}(f)) = f$ and $\lambda^{-1}(\lambda(b)) = b$ respectively. We recover ordinary function application via $\mathbf{app}(f, a) := \lambda^{-1}(f)[\mathbf{id}.a]$.

Although in practice our original formulation of function application is much more useful than anti- λ , the latter is more semantically natural. A variant of this argument is discussed by Gratzer et al. [Gra+22], because in the context of *modal type theories* one often encounters elimination forms akin to $\lambda^{-1}(-)$ and it can be far from obvious what the corresponding $\mathbf{app}(-, -)$ operation would be.

Exercise 2.15. Verify the claim that $\lambda^{-1}(-)$ and its β and η rules do in fact imply our original elimination, β , and η rules.

2.4.3 Dependent sums

We now present dependent pair types, also known as *dependent sums* or Σ -types. In a reversal of our discussion of Π -types, we will *begin* by defining dependent sums as an internalization of judgmental structure before unfolding this into inference rules.

The Σ type former behaves just like the Π type former: a natural family of types indexed by pairs of a type A and an A -indexed family of types B ,

$$\Sigma_{\Gamma} : (\sum_{A \in \text{Ty}(\Gamma)} \text{Ty}(\Gamma.A)) \rightarrow \text{Ty}(\Gamma)$$

or in inference rule notation,

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type}}{\Gamma \vdash \Sigma(A, B) \text{ type}} \quad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type}}{\Delta \vdash \Sigma(A, B)[\gamma] = \Sigma(A[\gamma], B[\gamma.A]) \text{ type}}$$

(Recall that we write $\sum_{A \in \text{Ty}(\Gamma)} \text{Ty}(\Gamma.A)$ for the indexed coproduct $\coprod_{A \in \text{Ty}(\Gamma)} \text{Ty}(\Gamma.A)$.)

Where Σ -types and Π -types differ is in their elements. Whereas $\Gamma \vdash \Pi(A, B)$ type internalizes terms with a free variable $\Gamma.A \vdash b : B$, the type $\Gamma \vdash \Sigma(A, B)$ type internalizes pairs of terms $\Gamma \vdash a : A$ and $\Gamma \vdash b : B[\text{id}.a]$, naturally in Γ :

$$\iota_{\Gamma, A, B} : \text{Tm}(\Gamma, \Sigma(A, B)) \cong \sum_{a \in \text{Tm}(\Gamma, A)} \text{Tm}(\Gamma, B[\text{id}.a])$$

Remarkably, the above line completes our definition of dependent sum types, but in the interest of the reader we will proceed to unfold this natural isomorphism into inference rules in three stages. First, we will unfold the maps $\iota_{\Gamma, A, B}$ and $\iota_{\Gamma, A, B}^{-1}$ into three term formers; second, we will unfold the two round-trip equations into a pair of equational rules; and finally, we will unfold the naturality condition into three more equational rules.

Remark 2.4.5. There is an unfortunate terminological collision between simple types and dependent types: although Π -types seem to generalize *simple functions*, they are called *dependent products*, and although Σ -types seem to generalize *simple products* because their elements are pairs, they are called *dependent sums*.

The reason is twofold: first, the elements of indexed coproducts (known to programmers as “tagged unions”) are actually pairs (“pairs of a tag bit with data”), whereas the elements of indexed products (“ n -ary pairs”) are actually functions (sending n to the n -th projection). Secondly, *both concepts* generalize simple finite products: the product $B_1 \times B_2$ is both an indexed product $\prod_{a \in \{1, 2\}} B_a$ and an indexed coproduct of a constant family $\sum_{_ \in B_1} B_2$. \diamond

To unpack the natural isomorphism, we note first that the forward direction $\iota_{\Gamma, A, B} : \text{Tm}(\Gamma, \Sigma(A, B)) \rightarrow \sum_{a \in \text{Tm}(\Gamma, A)} \text{Tm}(\Gamma, B[\text{id}.a])$ sends terms $\Gamma \vdash p : \Sigma(A, B)$ to (meta-)pairs of terms, so we can unfold this map into a pair of term formers with the same premises:

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash p : \Sigma(A, B)}{\Gamma \vdash \text{fst}(p) : A}$$

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash p : \Sigma(A, B)}{\Gamma \vdash \text{snd}(p) : B[\text{id}.\text{fst}(p)]}$$

The map $\iota_{\Gamma, A, B}^{-1} : \Sigma_{a \in \text{Term}(\Gamma, A)} \text{Term}(\Gamma, B[\text{id}.a]) \rightarrow \text{Term}(\Gamma, \Sigma(A, B))$ sends a pair of terms to a single term of type $\Sigma(A, B)$, so we unfold it into one term former with two term premises:

$$\frac{\Gamma \vdash a : A \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash b : B[\text{id}.a]}{\Gamma \vdash \mathbf{pair}(a, b) : \Sigma(A, B)}$$

Unlike in our judgmental analysis of dependent products, the standard introduction and elimination forms of dependent sums correspond exactly to the maps ι^{-1} and ι , so the two round-trip equations are exactly the standard β and η principles:

$$\frac{\Gamma \vdash a : A \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash b : B[\text{id}.a]}{\Gamma \vdash \mathbf{fst}(\mathbf{pair}(a, b)) = a : A \quad \Gamma \vdash \mathbf{snd}(\mathbf{pair}(a, b)) = b : B[\text{id}.a]}$$

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash p : \Sigma(A, B)}{\Gamma \vdash p = \mathbf{pair}(\mathbf{fst}(p), \mathbf{snd}(p)) : \Sigma(A, B)}$$

It remains to unpack the naturality of ι , which as we have seen previously, encodes the fact that the term formers commute with substitution. The reader may be surprised to learn, however, that the substitution rule for $\mathbf{pair}(-, -)$ actually implies the substitution rules for $\mathbf{fst}(-)$ and $\mathbf{snd}(-)$ in the presence of β and η . (Categorically, this is the fact that naturality of ι^{-1} implies naturality of ι , as we saw in Exercise 2.12.) Given the rule

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : A \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash b : B[\text{id}.a]}{\Delta \vdash \mathbf{pair}(a, b)[\gamma] = \mathbf{pair}(a[\gamma], b[\gamma]) : \Sigma(A, B)[\gamma]}$$

fix a substitution $\Delta \vdash \gamma : \Gamma$ and a term $\Gamma \vdash p : \Sigma(A, B)$. Then

$$\begin{aligned} & \mathbf{fst}(p)[\gamma] \\ &= \mathbf{fst}(\mathbf{pair}(\mathbf{fst}(p)[\gamma], \mathbf{snd}(p)[\gamma])) && \text{by the } \beta \text{ rule} \\ &= \mathbf{fst}(\mathbf{pair}(\mathbf{fst}(p), \mathbf{snd}(p))[\gamma]) && \text{by the above rule} \\ &= \mathbf{fst}(p[\gamma]) && \text{by the } \eta \text{ rule} \end{aligned}$$

and the calculation for $\mathbf{snd}(-)$ is identical. Nevertheless it is typical to include substitution rules for all three term formers: there is nothing wrong with equating terms that are already equal, and even in type theory, discretion can be the better part of valor.

Exercise 2.16. Check that the substitution rule for \mathbf{pair} above is meta-well-typed, in particular the second component $b[\gamma]$. (Hint: use Exercise 2.3.)

Exercise 2.17. Show that the substitution rule for $\lambda^{-1}(-)$ follows from the substitution rule for $\lambda(-)$ and the equations $\lambda(\lambda^{-1}(f)) = f$ and $\lambda^{-1}(\lambda(b)) = b$.

2.4.4 Extensional equality

We now turn to the simplest form of propositional equality, known as *extensional equality* or *Eq-types*. As their name suggests, **Eq**-types internalize the term equality judgment. They are defined as follows, naturally in Γ :

$$\mathbf{Eq}_\Gamma : (\sum_{A \in \mathbf{Ty}(\Gamma)} \mathbf{Tm}(\Gamma, A) \times \mathbf{Tm}(\Gamma, A)) \rightarrow \mathbf{Ty}(\Gamma)$$

$$\iota_{\Gamma, A, a, b} : \mathbf{Tm}(\Gamma, \mathbf{Eq}(A, a, b)) \cong \{\star \mid a = b\}$$

In other words, $\mathbf{Eq}(A, a, b)$ is a type when $\Gamma \vdash a : A$ and $\Gamma \vdash b : A$, and has a unique inhabitant exactly when the judgment $\Gamma \vdash a = b : A$ holds (otherwise it is empty). The inference rules for extensional equality are as follows:

$$\frac{\Gamma \vdash a, b : A}{\Gamma \vdash \mathbf{Eq}(A, a, b) \text{ type}} \qquad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a, b : A}{\Delta \vdash \mathbf{Eq}(A, a, b)[\gamma] = \mathbf{Eq}(A[\gamma], a[\gamma], b[\gamma]) \text{ type}}$$

$$\frac{\Gamma \vdash a : A}{\Gamma \vdash \mathbf{refl} : \mathbf{Eq}(A, a, a)} \qquad \frac{\Gamma \vdash a, b : A \quad \Gamma \vdash p : \mathbf{Eq}(A, a, b)}{\Gamma \vdash a = b : A}$$

$$\frac{\Gamma \vdash a, b : A \quad \Gamma \vdash p : \mathbf{Eq}(A, a, b)}{\Gamma \vdash p = \mathbf{refl} : \mathbf{Eq}(A, a, b)}$$

The penultimate rule is known as *equality reflection*, and it is somewhat unusual because it concludes an arbitrary term equality judgment from the existence of a term. This rule is quite strong in light of the facts that (1) judgmentally equal terms can be silently exchanged at any location in any judgment, (2) the equality proof $\Gamma \vdash p : \mathbf{Eq}(A, a, b)$ is not recorded in those exchanges, and (3) p could even be a variable, e.g., in context $\Gamma.\mathbf{Eq}(A, a, b)$.

Type theories with an extensional equality type are called *extensional*. The consequences of equality reflection will be the primary motivation behind the latter half of this book, but for now we simply note that these rules are a very natural axiomatization of an equality type as the internalization of equality.

Exercise 2.18. Explain how these inference rules correspond to our \mathbf{Eq}_Γ and $\iota_{\Gamma, A, a, b}$ definition.

Exercise 2.19. Where are the substitution rules for term formers? (Hint: there are two equivalent answers, in terms of either the natural isomorphism or the inference rules.)

2.4.5 The unit type

We conclude our tour of the best-behaved connectives of type theory with the simplest connective of all: the unit type.

$$\begin{aligned} \mathbf{Unit}_\Gamma &\in \mathbf{Ty}(\Gamma) \\ \iota_\Gamma : \mathbf{Tm}(\Gamma, \mathbf{Unit}) &\cong \{\star\} \end{aligned}$$

This unfolds to the following rules:

$$\begin{array}{c} \frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{Unit} \text{ type}} \qquad \frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{Unit}[\gamma] = \mathbf{Unit} \text{ type}} \\ \frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{tt} : \mathbf{Unit}} \qquad \frac{\Gamma \vdash a : \mathbf{Unit}}{\Gamma \vdash a = \mathbf{tt} : \mathbf{Unit}} \end{array}$$

Exercise 2.20. Where is the elimination principle? Where are the substitution rules for term formers? (Hint: what would these say in terms of the natural isomorphism?)

2.5 Inductive types: Void, Bool, Nat

We now turn our attention to *inductive types*, data types with induction principles. Unlike the type formers in Section 2.4, which are typically “hard coded” into type theories,⁸ inductive types are usually specified by users as extensions to the theory via inductive schemas [Dyb94; CP90a] (essentially, data type declarations), or in theoretical contexts, encoded as well-founded trees known as **W**-types [ML82; ML84b]. These schemas can be extended *ad infinitum* to account for increasingly complex forms of inductive definition, including indexed induction [Dyb94], mutual induction, induction-recursion [Dyb00], induction-induction [NFS12], quotient induction-induction [KKA19], and so forth.

For simplicity we restrict our attention to three examples—the empty type, booleans, and natural numbers—that illustrate the basic issues that arise when specifying inductive types in type theory. Unfortunately, we will immediately need to refine Slogan 2.4.4.

⁸This is an oversimplification: in practice, Σ and \mathbf{Unit} are usually obtained as special cases of *dependent record types* [Pol02], n -ary Σ -types with named projections.

2.5.1 The empty type

We begin with the empty type **Void**, a “type with no elements.” Logically, this type corresponds to the false proposition, so there should be no way to construct an element of **Void** (a proof of false) except by deriving a contradiction from local hypotheses. The type former is straightforward: naturally in Γ , a constant $\mathbf{Void}_\Gamma \in \text{Ty}(\Gamma)$, or

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{Void} \text{ type}} \qquad \frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{Void}[\gamma] = \mathbf{Void} \text{ type}}$$

As for the elements of **Void**, an obvious guess is to say that the elements of the empty type at each context are the empty set, i.e., naturally in Γ ,

$$t_\Gamma : \text{Tm}(\Gamma, \mathbf{Void}) \cong \emptyset \qquad (!?)$$

This cannot be right, however, because **Void** *does* have elements in some contexts—the variable rule alone forces $\mathbf{q} \in \text{Tm}(\Gamma.\mathbf{Void}, \mathbf{Void})$, and other type formers can populate **Void** even further, e.g., $\mathbf{app}(\mathbf{q}, \mathbf{tt}) \in \text{Tm}(\Gamma.\Pi(\mathbf{Unit}, \mathbf{Void}), \mathbf{Void})$.

Interlude: mapping in, mapping out To see how to proceed, let us take a brief sojourn into set theory. There are several ways to define the product $A \times B$ of two sets, for example by constructing it as the set of ordered pairs $\{(a, b) \mid a \in A \wedge b \in B\}$ or even more explicitly as the set $\{\{\{a\}, \{a, b\}\} \mid a \in A \wedge b \in B\}$. However, in addition to these explicit constructions, it is also possible to *characterize* the set $A \times B$ up to isomorphism, as the set such that every function $X \rightarrow A \times B$ is determined by a pair of functions $X \rightarrow A$ and $X \rightarrow B$ and vice versa.

Similarly, we can characterize one-element sets **1** as those sets for which there is exactly one function $X \rightarrow \mathbf{1}$ for all sets X . In fact, both of these characterizations are set-theoretical analogues of Slogan 2.4.4, where X plays the role of the context Γ .

After some thought, we realize that the analogous characterization of the zero-element (empty) set **0** is significantly more awkward: there is exactly one function $X \rightarrow \mathbf{0}$ when X is itself empty, and no functions $X \rightarrow \mathbf{0}$ when X is non-empty. As it turns out, in this case it is more elegant to consider the functions *out* of **0** rather than the functions *into* it: a zero-element set **0** has exactly one function $\mathbf{0} \rightarrow X$ for all sets X .

Exercise 2.21. Suppose that Z is a set such that for all sets X there is exactly one function $Z \rightarrow X$. Show that Z is isomorphic to the empty set.

Void revisited Recall from Section 2.3 that terms correspond to “dependent functions from Γ to A .” In Section 2.4 we considered only type formers T that are easily characterized in terms of the maps *into* that type former from an arbitrary context Γ : in each case we defined maps/terms $\text{Tm}(\Gamma, T)$ as naturally isomorphic to the data of T ’s introduction rule.

To characterize the maps *out of* **Void** into an arbitrary type A , we cannot leave the context fully unconstrained; instead, we must characterize the maps/terms $\text{Tm}(\Gamma.\mathbf{Void}, A)$ for all $\vdash \Gamma \text{ cx}$ and $\Gamma.\mathbf{Void} \vdash A$ type, recalling that—by the rules for Π -types—these are equivalently the dependent functions out of **Void** in context Γ , i.e., $\Gamma \vdash f : \Pi(\mathbf{Void}, A)$.

Advanced Remark 2.5.1. Writing \mathcal{C} for the category of contexts and substitutions, terms $\text{Tm}(\Gamma, A)$ are indeed “dependent morphisms” from Γ to A ; more precisely, by Exercise 2.2, they are ordinary morphisms $\Gamma \rightarrow \Gamma.A$ in the slice category \mathcal{C}/Γ . Thus, for *right adjoint* type operations G —those in Section 2.4—it is easy to describe $\text{Tm}(\Gamma, G(A))$ directly.

For *left adjoint* type operations F , the situation is more fraught. Type theory is fundamentally “right-biased” because its judgments concern maps from arbitrary contexts *into* fixed types, but not vice versa. Thus to discuss dependent morphisms $F(X) \rightarrow A$ we must speak about elements of $\text{Tm}(\Gamma.F(X), A)$, quantifying not only over the ambient context/slice Γ but also the type A into which we are mapping.

Confusingly, we encountered no issues defining Σ -types, despite dependent sum being the left adjoint to pullback. This is because Σ is also the right adjoint to the functor $\mathcal{C} \rightarrow \mathcal{C}^{\rightarrow}$ sending $A \mapsto \mathbf{id}_A$, and it is the latter perspective that we axiomatize. The left adjoint axiomatization makes an appearance in some systems—particularly in the context of programming languages with existential types—phrased as **let** $(a, b) = p$ in x . \diamond

Putting all these ideas together, we will define **Void** as the type for which, naturally in Γ , there is exactly one dependent function from **Void** to A for any dependent type A :

$$\rho_{\Gamma, A} : \text{Tm}(\Gamma.\mathbf{Void}, A) \cong \{\star\}$$

To sum up the difference between the incorrect definition $\text{Tm}(\Gamma, \mathbf{Void}) \cong \emptyset$ and the correct one above, the former states that $\text{Tm}(\Gamma, \mathbf{Void})$ is the smallest set (in the sense of mapping into all other sets), whereas the latter states that in any context, **Void** is the smallest *type*. More poetically, at the level of judgments we can see that **Void** is not always empty, but at the level of types, every type “believes” that **Void** is empty.

Unwinding $\rho_{\Gamma, A}$ into inference rules, we obtain:

$$\frac{\vdash \Gamma \text{ cx} \quad \Gamma.\mathbf{Void} \vdash A \text{ type}}{\Gamma.\mathbf{Void} \vdash \mathbf{absurd}' : A} \quad \frac{\vdash \Gamma \text{ cx} \quad \Gamma.\mathbf{Void} \vdash a : A}{\Gamma.\mathbf{Void} \vdash \mathbf{absurd}' = a : A}$$

We have marked these rules with 📎 to indicate that they are provisional; in practice, as we previously discussed for $\lambda^{-1}(-)$, it is awkward to use rules whose conclusions constrain the shape of their context. But just as with **app**($-$, $-$), it is more standard to present an

equivalent axiomatization $\mathbf{absurd}(b) := \mathbf{absurd}'[\mathbf{id}.b]$ that “builds in a cut”:

$$\frac{\Gamma \vdash b : \mathbf{Void} \quad \Gamma.\mathbf{Void} \vdash A \text{ type}}{\Gamma \vdash \mathbf{absurd}(b) : A[\mathbf{id}.b]} \quad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash b : \mathbf{Void} \quad \Gamma.\mathbf{Void} \vdash A \text{ type}}{\Delta \vdash \mathbf{absurd}(b)[\gamma] = \mathbf{absurd}(b[\gamma]) : A[\gamma.b[\gamma]]}$$

$$\frac{\Gamma \vdash b : \mathbf{Void} \quad \Gamma.\mathbf{Void} \vdash a : A}{\Gamma \vdash \mathbf{absurd}(b) = a[\mathbf{id}.b] : A[\mathbf{id}.b]} \text{ } \textcircled{e}$$

The term $\mathbf{absurd}(-)$ is known as the *induction principle* for \mathbf{Void} , in the sense that it allows users to prove a theorem for all terms of type \mathbf{Void} by proving that it holds for each constructor of \mathbf{Void} , of which there are none.

In light of our definition of \mathbf{Void} , we update Slogan 2.4.4 as follows:

Slogan 2.5.2. *A connective in type theory is given by (1) a natural type-forming operation Υ and (2) one of the following:*

- 2.1. *a natural isomorphism relating $\mathbf{Tm}(\Gamma, \Upsilon)$ to judgmentally-determined structure, or*
- 2.2. *for all $\Gamma.\Upsilon \vdash A \text{ type}$, a natural isomorphism relating $\mathbf{Tm}(\Gamma.\Upsilon, A)$ to judgmentally-determined structure.*

The final rule for $\mathbf{absurd}(-)$, the η principle, implies a very strong equality principle for terms in an inconsistent context (Exercise 2.25) which we derive in the following sequence of exercises. For this reason, and because this rule is derivable in the presence of extensional equality (Section 2.5.4), we consider it provisional \textcircled{e} for the time being.

Exercise 2.22. Show that if $\Gamma \vdash b_0, b_1 : \mathbf{Void}$ then $\Gamma \vdash b_0 = b_1 : \mathbf{Void}$.

Exercise 2.23. Fixing $\Delta \vdash \gamma : \Gamma$, prove that there is at most one substitution $\Delta \vdash \bar{\gamma} : \Gamma.\mathbf{Void}$ satisfying $\mathbf{p} \circ \bar{\gamma} = \gamma$.

Exercise 2.24. Let $\Gamma.\mathbf{Void} \vdash A \text{ type}$ and $\Gamma \vdash a : A[\mathbf{id}.b]$. Show that $\Gamma.\mathbf{Void} \vdash A[\mathbf{id}.b \circ \mathbf{p}] = A \text{ type}$, and therefore that $\Gamma.\mathbf{Void} \vdash a[\mathbf{p}] : A$.

Exercise 2.25. Derive the following rule, using the previous exercise as well as the η rule.

$$\frac{\Gamma \vdash b : \mathbf{Void} \quad \Gamma.\mathbf{Void} \vdash A \text{ type} \quad \Gamma \vdash a : A[\mathbf{id}.b]}{\Gamma \vdash a = \mathbf{absurd}(b) : A[\mathbf{id}.b]} \Rightarrow$$

Exercise 2.26. We have included the rule $\Delta \vdash \mathbf{absurd}(b)[\gamma] = \mathbf{absurd}(b[\gamma]) : A[\gamma.b[\gamma]]$ but it is in fact derivable using the η rule. Prove this.

2.5.2 Booleans

We turn now to the booleans **Bool**, a “type with two elements.” Once again the type former is straightforward: $\mathbf{Bool}_\Gamma \in \text{Ty}(\Gamma)$ naturally in Γ , or

$$\frac{}{\Gamma \vdash \mathbf{Bool} \text{ type}} \qquad \frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{Bool}[\gamma] = \mathbf{Bool} \text{ type}}$$

It is also clear that we want two constructors of **Bool**, **true** and **false**, natural in Γ :

$$\frac{}{\Gamma \vdash \mathbf{true} : \mathbf{Bool}} \qquad \frac{}{\Gamma \vdash \mathbf{false} : \mathbf{Bool}}$$

$$\frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{true} = \mathbf{true}[\gamma] : \mathbf{Bool}} \qquad \frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{false} = \mathbf{false}[\gamma] : \mathbf{Bool}}$$

Keeping Slogan 2.5.2 in mind, there are two possible ways for us to complete our axiomatization of **Bool**. As with **Void** it is tempting but incorrect to define $\iota : \text{Tm}(\Gamma, \mathbf{Bool}) \cong \{\star, \star'\}$; although the natural transformation ι^{-1} is equivalent to our rules for **true** and **false**, ι does not account for variables of type **Bool** or other indeterminate booleans that arise in non-empty contexts.⁹ Thus we must instead characterize maps *out of* **Bool** by giving a family of sets naturally isomorphic to $\text{Tm}(\Gamma.\mathbf{Bool}, A)$.

So, what should terms $\Gamma.\mathbf{Bool} \vdash a : A$ be? By substitution, such a term clearly determines a pair of terms $\Gamma \vdash a[\mathbf{id.true}] : A[\mathbf{id.true}]$ and $\Gamma \vdash a[\mathbf{id.false}] : A[\mathbf{id.false}]$. Conversely, if **true** and **false** are the “only” booleans, then such a pair of terms should uniquely determine elements of $\text{Tm}(\Gamma.\mathbf{Bool}, A)$ in the sense that to map out of **Bool**, it suffices to explain what to do on **true** and on **false**.

To formalize this idea, let us write $((\mathbf{id.true})^*, (\mathbf{id.false})^*)$ for the function which sends $a \in \text{Tm}(\Gamma.\mathbf{Bool}, A)$ to the pair $(a[\mathbf{id.true}], a[\mathbf{id.false}])$. We complete our specification of **Bool** by asking for this map to be a natural isomorphism; thus, naturally in Γ , we have:

$$\mathbf{Bool}_\Gamma \in \text{Ty}(\Gamma)$$

$$\mathbf{true}_\Gamma, \mathbf{false}_\Gamma \in \text{Tm}(\Gamma, \mathbf{Bool})$$

$$((\mathbf{id.true})^*, (\mathbf{id.false})^*) : \text{Tm}(\Gamma.\mathbf{Bool}, A) \cong \text{Tm}(\Gamma, A[\mathbf{id.true}]) \times \text{Tm}(\Gamma, A[\mathbf{id.false}])$$

This definition is remarkable in several ways. For the first time we are asking not only for the existence of some natural isomorphism, but for a *particular map* to be a natural isomorphism; and because this map is defined in terms of **true** and **false**, these must be asserted prior to the natural isomorphism itself. We update our slogan accordingly:

⁹Even if variables $x : \mathbf{Bool}$ stand for one of **true** or **false**, x itself must be an indeterminate boolean equal to neither constructor; otherwise the identity $\lambda x.x : \mathbf{Bool} \rightarrow \mathbf{Bool}$ would be a constant function.

Slogan 2.5.3. A connective in type theory is given by (1) a natural type-forming operation Υ and (2) one of the following:

- 2.1. a natural isomorphism relating $\text{Tm}(\Gamma, \Upsilon)$ to judgmentally-determined structure, or
- 2.2. a collection of natural term constructors for Υ which, for all $\Gamma. \Upsilon \vdash A$ type, determine a natural isomorphism relating $\text{Tm}(\Gamma, \Upsilon, A)$ to judgmentally-determined structure.

In the case of **Void** we simply had no term constructors to specify, and because there is at most one (natural) isomorphism between anything and $\{\star\}$, it was unnecessary for us to specify the underlying map. In general, however, we emphasize that it is essential to specify the map; this is what ensures that when we define a function “by cases” on **true** and **false**, applying it to **true** or **false** recovers the specified case and not something else. On the other hand, because we have specified the underlying map, it being an isomorphism is a *property* rather than additional structure: there is at most one possible inverse.

Zooming out, however, our definition of **Bool** has a similar effect to our definition of **Void** from Section 2.5.1: $\text{Tm}(\Gamma, \mathbf{Bool})$ is *not* the set $\{\mathbf{true}, \mathbf{false}\}$ at the level of judgments, but every type “believes” that it is. This is the role of type-theoretic induction principles.

Advanced Remark 2.5.4. From the categorical perspective, option 2.2 in Slogan 2.5.3 asserts that the inclusion map of Υ ’s constructors into Υ ’s terms is *left orthogonal* to all types. Maps which are left orthogonal to a class of objects and whose codomain belongs to that class are known as *fibrant replacements*; in this sense, we have defined $\text{Tm}(-, \mathbf{Void})$ and $\text{Tm}(-, \mathbf{Bool})$ as fibrant replacements of the constantly zero- and two-element presheaves. This perspective is crucial to early work in homotopy type theory [AW09] and the formulation of the intensional identity type in natural models [Awo18]. \diamond

It remains to unfold our natural isomorphism into inference rules. We do not need any additional rules for the forward map, which is substitution by **id.true** and **id.false**. As the reader may have already guessed, the backward map is essentially¹⁰ dependent **if**:

$$\frac{\Gamma.\mathbf{Bool} \vdash A \text{ type} \quad \Gamma \vdash a_t : A[\mathbf{id.true}] \quad \Gamma \vdash a_f : A[\mathbf{id.false}] \quad \Gamma \vdash b : \mathbf{Bool}}{\Gamma \vdash \mathbf{if}(a_t, a_f, b) : A[\mathbf{id.b}]}$$

$$\frac{\Gamma.\mathbf{Bool} \vdash A \text{ type} \quad \Gamma \vdash a_t : A[\mathbf{id.true}] \quad \Gamma \vdash a_f : A[\mathbf{id.false}] \quad \Gamma \vdash b : \mathbf{Bool} \quad \Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{if}(a_t, a_f, b)[\gamma] = \mathbf{if}(a_t[\gamma], a_f[\gamma], b[\gamma]) : A[\gamma.b[\gamma]}}$$

¹⁰The inverse directly lands in $\Gamma.\mathbf{Bool}$ and not Γ , but as with **absurd**’ (Section 2.5.1) we adopt a more standard presentation in which all conclusions have a generic context; see Exercise 2.27.

The fact that `if` is an inverse to $((\mathbf{id.true})^*, (\mathbf{id.false})^*)$ expresses the β and η laws:

$$\frac{\Gamma.\mathbf{Bool} \vdash A \text{ type} \quad \Gamma \vdash a_t : A[\mathbf{id.true}] \quad \Gamma \vdash a_f : A[\mathbf{id.false}]}{\Gamma \vdash \mathbf{if}(a_t, a_f, \mathbf{true}) = a_t : A[\mathbf{id.true}] \quad \Gamma \vdash \mathbf{if}(a_t, a_f, \mathbf{false}) = a_f : A[\mathbf{id.false}]}$$

$$\frac{\Gamma.\mathbf{Bool} \vdash A \text{ type} \quad \Gamma.\mathbf{Bool} \vdash a : A \quad \Gamma \vdash b : \mathbf{Bool}}{\Gamma \vdash \mathbf{if}(a[\mathbf{id.true}], a[\mathbf{id.false}], b) = a[\mathbf{id.b}] : A[\mathbf{id.b}]}$$

The β laws—the first two equations—are perhaps more familiar than the η law, which effectively asserts that any two terms dependent on **Bool** are equal if (and only if) they are equal on **true** and **false**. (The η rule is sometimes decomposed into a “local expansion” and a collection of “commuting conversions.”) Although semantically justified, it is typical to omit judgmental η laws for all inductive types because they are not syntax-directed and thus challenging to implement, and because they are derivable in the presence of extensional equality (Section 2.5.4).

Exercise 2.27. Give rules axiomatizing the boolean analogue of **absurd'**, and prove that these rules are interderivable with our rules for $\mathbf{if}(a_t, a_f, b)$.

2.5.3 Natural numbers

Our final example of an inductive type is the type of natural numbers **Nat**, the “least type closed under $\mathbf{zero} : \mathbf{Nat}$ and $\mathbf{suc}(-) : \mathbf{Nat} \rightarrow \mathbf{Nat}$.” The natural numbers more or less fit the same pattern as **Void** and **Bool**, but the recursive nature of $\mathbf{suc}(-)$ complicates the situation significantly. The formation and introduction rules remain straightforward:

$$\frac{}{\Gamma \vdash \mathbf{Nat} \text{ type}} \quad \frac{}{\Gamma \vdash \mathbf{zero} : \mathbf{Nat}} \quad \frac{\Gamma \vdash n : \mathbf{Nat}}{\Gamma \vdash \mathbf{suc}(n) : \mathbf{Nat}}$$

$$\frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{Nat}[\gamma] = \mathbf{Nat} \text{ type}} \quad \frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{zero}[\gamma] = \mathbf{zero} : \mathbf{Nat}}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash n : \mathbf{Nat}}{\Gamma \vdash \mathbf{suc}(n)[\gamma] = \mathbf{suc}(n[\gamma]) : \mathbf{Nat}}$$

Following the pattern we established with **Bool**, we might ask for maps out of **Nat** to be determined by their behavior on \mathbf{zero} and $\mathbf{suc}(-)$, i.e., for the two substitutions

$$(\mathbf{id.zero})^* : \mathbf{Tm}(\Gamma.\mathbf{Nat}, A) \rightarrow \mathbf{Tm}(\Gamma, A[\mathbf{id.zero}])$$

$$(\mathbf{p.suc}(q))^* : \mathbf{Tm}(\Gamma.\mathbf{Nat}, A) \rightarrow \mathbf{Tm}(\Gamma.\mathbf{Nat}, A[\mathbf{p.suc}(q)])$$

to determine, for every $\Gamma.\text{Nat} \vdash A$ type, a natural isomorphism

$$\begin{aligned} & ((\text{id.zero})^*, (\text{p.suc(q)})^*) : \\ & \text{Tm}(\Gamma.\text{Nat}, A) \cong \text{Tm}(\Gamma, A[\text{id.zero}]) \times \text{Tm}(\Gamma.\text{Nat}, A[\text{p.suc(q)}]) \quad (!?) \end{aligned}$$

This turns out to not be the correct definition, but first, note that the first substitution moves us from $\Gamma.\text{Nat}$ to Γ (analogously to **Bool**) whereas the second substitution moves us from $\Gamma.\text{Nat}$ also to $\Gamma.\text{Nat}$; this is because the $\text{suc}(-)$ constructor has type “ $\text{Nat} \rightarrow \text{Nat}$,” so the condition of “being determined by one’s behavior on $\text{suc}(n) : \text{Nat}$ ” is properly stated relative to a variable $n : \text{Nat}$. Put more simply, if the argument of $\text{suc}(-)$ was of type X rather than Nat , then the latter substitution would be $\Gamma.X \vdash \text{p.suc(q)} : \Gamma.\text{Nat}$.

But given that $\text{suc}(-)$ is recursive—taking Nat to Nat —we now for the first time are defining a judgment by a natural isomorphism whose *right-hand side also* has the very same judgment we are trying to define, namely $\text{Tm}(\Gamma.\text{Nat}, \dots)$, i.e., terms in context $\Gamma.\text{Nat}$. This natural isomorphism is therefore not so much a *definition* of its left-hand side as it is an *equation* that the left-hand side must satisfy—in principle, this equation may have many different solutions for $\text{Tm}(\Gamma.\text{Nat}, A)$, or no solutions at all.

Interlude: initial algebras This equation asserts in essence that the natural numbers are a set N satisfying the isomorphism $N \cong \{\star\} + N$,¹¹ where the reverse map equips N with a choice of “implementations” of $\text{zero} \in N$ and $\text{suc}(-) : N \rightarrow N$. The set of natural numbers \mathbb{N} with $\text{zero} := 0$ and $\text{suc}(n) := n + 1$ are a solution, but there are infinitely many *other* solutions as well, such as $\mathbb{N} + \{\infty\}$ with $\text{zero} := 0$, $\text{suc}(n) := n + 1$, and $\text{suc}(\infty) := \infty$.

Nevertheless one might imagine that $(\mathbb{N}, 0, - + 1)$ is a distinguished solution in some way, and indeed it is the “least” set N with a point $z \in N$ and endofunction $s : N \rightarrow N$ —here we are dropping the requirement of (z, s) being an isomorphism—in the sense that for any (N, z, s) there is a unique function $f : \mathbb{N} \rightarrow N$ with $f(0) = z$ and $f(n + 1) = s(f(n))$. Such triples (N, z, s) are known as *algebras* for the signature $N \mapsto 1 + N$, structure-preserving functions between algebras are known as *algebra homomorphisms*, and algebras with the above minimality property are *initial algebras*.

The above definitions extend straightforwardly to dependent algebras and homomorphisms: given an ordinary algebra (N, z, s) , a *displayed algebra over* (N, z, s) is a triple of an N -indexed family of sets $\{\tilde{N}_n\}_{n \in N}$, an element $\tilde{z} \in \tilde{N}_z$, and a function $\tilde{s} : (n : N) \rightarrow \tilde{N}_n \rightarrow \tilde{N}_{s(n)}$ [KKA19]. Given any displayed algebra $(\tilde{N}, \tilde{z}, \tilde{s})$ over the natural number algebra $(\mathbb{N}, 0, - + 1)$, there is once again a unique function $f : (n : \mathbb{N}) \rightarrow \tilde{N}_n$ with $f(0) = \tilde{z}$ and $f(n + 1) = \tilde{s}(n, f(n))$. The reader is likely familiar with the special case of displayed algebras over \mathbb{N} valued in *propositions* rather than sets:

$$\forall P : \mathbb{N} \rightarrow \mathbf{Prop}. P(0) \implies (\forall n. P(n) \implies P(n + 1)) \implies \forall n. P(n)$$

¹¹Why? In algebraic notation and ignoring dependency, the equation states that $A^{\Gamma \times N} \cong A^\Gamma \times A^{\Gamma \times N}$, which simplifies to $(\Gamma \times N) \cong \Gamma + (\Gamma \times N)$ and thus $N \cong 1 + N$.

Advanced Remark 2.5.5. The data of a displayed algebra over (N, z, s) is equivalent to the data of an algebra homomorphism into (N, z, s) , where the forward direction of this equivalence sends the family $\{\tilde{N}_n\}_{n \in \mathbb{N}}$ to the first projection $(\sum_{n \in \mathbb{N}} \tilde{N}_n) \rightarrow N$. A displayed algebra over the natural number algebra is thus a homomorphism $\tilde{N} \rightarrow \mathbb{N}$; the initiality of \mathbb{N} implies this map has a unique section homomorphism, which unfolds to the dependent universal property stated above. \diamond

Natural numbers revisited Coming back to our specification of **Nat**, our formation and introduction rules axiomatize an algebra $(\mathbf{Nat}, \mathbf{zero}, \mathbf{suc}(-))$ for the signature $N \mapsto 1 + N$, but our proposed **Bool**-style natural isomorphism does not imply that this algebra is initial. The solution is to simply axiomatize that any displayed algebra over $(\mathbf{Nat}, \mathbf{zero}, \mathbf{suc}(-))$ admits a unique displayed algebra homomorphism from $(\mathbf{Nat}, \mathbf{zero}, \mathbf{suc}(-))$.

Unwinding definitions, we ask that naturally in Γ , and for any $A \in \text{Ty}(\Gamma.\mathbf{Nat})$, $a_z \in \text{Tm}(\Gamma, A[\mathbf{id.zero}])$, and $a_s \in \text{Tm}(\Gamma.\mathbf{Nat}.A, A[\mathbf{p}^2.\mathbf{suc}(\mathbf{q}[\mathbf{p}])])$, we have an isomorphism:

$$\rho_{\Gamma, A, a_z, a_s} : \{a \in \text{Tm}(\Gamma.\mathbf{Nat}, A) \mid a_z = a[\mathbf{id.zero}] \wedge a_s[\mathbf{p.q.a}] = a[\mathbf{p.suc}(\mathbf{q})]\} \cong \{\star\}$$

The type of a_s is easier to understand with named variables: it is a term of type $A(\mathbf{suc}(n))$ in context $\Gamma, n : \mathbf{Nat}, a : A(n)$.

Remark 2.5.6. This is the third time we have defined a connective in terms of a natural isomorphism with $\{\star\}$. In Section 2.4.5, we used such an isomorphism to assert that **Unit** has a unique element in every context; in Section 2.5.1, we asserted dually that every dependent type over **Void** admits a unique dependent function from **Void**. The present definition is analogous to the latter, but restricted to algebras: every displayed algebra over **Nat** admits a unique displayed algebra homomorphism from **Nat**. \diamond

Advanced Remark 2.5.7. In light of Remark 2.5.4 and Remark 2.5.6, we have defined **Nat** as the fibrant replacement of the initial object in the category of $(1 + -)$ -algebras. \diamond

In rule form, the reverse direction of the natural isomorphism states that any displayed algebra (A, a_z, a_s) over **Nat** gives rise to a map out of **Nat**,

$$\frac{\Gamma.\mathbf{Nat} \vdash A \text{ type} \quad \Gamma \vdash n : \mathbf{Nat} \quad \Gamma \vdash a_z : A[\mathbf{id.zero}] \quad \Gamma.\mathbf{Nat}.A \vdash a_s : A[\mathbf{p}^2.\mathbf{suc}(\mathbf{q}[\mathbf{p}])]}{\Gamma \vdash \mathbf{rec}(a_z, a_s, n) : A[\mathbf{id}.n]}$$

which commutes with substitution,

$$\frac{\Gamma.\mathbf{Nat} \vdash A \text{ type} \quad \Delta \vdash \gamma : \Gamma \quad \Gamma \vdash n : \mathbf{Nat} \quad \Gamma \vdash a_z : A[\mathbf{id.zero}] \quad \Gamma.\mathbf{Nat}.A \vdash a_s : A[\mathbf{p}^2.\mathbf{suc}(\mathbf{q}[\mathbf{p}])]}{\Delta \vdash \mathbf{rec}(a_z, a_s, n)[\gamma] = \mathbf{rec}(a_z[\gamma], a_s[\gamma.\mathbf{Nat}.A], n[\gamma]) : A[\gamma.n[\gamma]]}$$

and is a displayed algebra homomorphism, i.e., sends **zero** to a_z and **suc**(n) to $a_s(n, \mathbf{rec}(n))$:

$$\frac{\Gamma.\mathbf{Nat} \vdash A \text{ type} \quad \Gamma \vdash a_z : A[\mathbf{id.zero}] \quad \Gamma.\mathbf{Nat}.A \vdash a_s : A[\mathbf{p}^2.\mathbf{suc}(\mathbf{q}[\mathbf{p}])]}{\Gamma \vdash \mathbf{rec}(a_z, a_s, \mathbf{zero}) = a_z : A[\mathbf{id.zero}]}$$

$$\frac{\Gamma \vdash n : \mathbf{Nat} \quad \Gamma.\mathbf{Nat} \vdash A \text{ type} \quad \Gamma \vdash a_z : A[\mathbf{id.zero}] \quad \Gamma.\mathbf{Nat}.A \vdash a_s : A[\mathbf{p}^2.\mathbf{suc}(\mathbf{q}[\mathbf{p}])]}{\Gamma \vdash \mathbf{rec}(a_z, a_s, \mathbf{suc}(n)) = a_s[\mathbf{id.n.rec}(a_z, a_s, n)] : A[\mathbf{id.suc}(n)]}$$

Finally, the η rule of **Nat**, which is again typically omitted, expresses that there is exactly one displayed algebra homomorphism from **Nat** to (A, a_z, a_s) : if $\Gamma.\mathbf{Nat} \vdash a : A$ is a term that sends **zero** to a_z and **suc**(n) to $a_s(n, a[\mathbf{id.n}])$, then it is equal to $\mathbf{rec}(a_z, a_s, \mathbf{q})$.

$$\frac{\Gamma.\mathbf{Nat} \vdash A \text{ type} \quad \Gamma.\mathbf{Nat} \vdash a : A \quad \Gamma \vdash n : \mathbf{Nat} \quad \Gamma \vdash a_z : A[\mathbf{id.zero}] \quad \Gamma \vdash a_z = a[\mathbf{id.zero}] : A[\mathbf{id.zero}] \quad \Gamma.\mathbf{Nat}.A \vdash a_s : A[\mathbf{p}^2.\mathbf{suc}(\mathbf{q}[\mathbf{p}])]}{\Gamma.\mathbf{Nat} \vdash a_s[\mathbf{p.q.a}] = a[\mathbf{p.suc}(\mathbf{q})] : A[\mathbf{p.suc}(\mathbf{q})]} \quad \Gamma \vdash \mathbf{rec}(a_z, a_s, n) = a[\mathbf{id.n}] : A[\mathbf{id.n}]$$

Exercise 2.28. Rewrite the first **rec** rule using named variables instead of **p** and **q**, and convince yourself that it expresses a form of natural number induction.

Exercise 2.29. Define addition for **Nat** in terms of **rec**. We strongly recommend solving Exercise 2.28 prior to this exercise in order to use standard named syntax.

Inductive types are initial algebras Our definition of **Nat** is more similar to our definitions of **Void** and **Bool** than it may first appear. In fact, all three types are initial algebras for different signatures, although the absence of recursive constructors in **Void** and **Bool** allowed us to sidestep this machinery. The empty type **Void** is the initial algebra for the signature $X \mapsto \mathbf{0}$: a (displayed) **0**-algebra is just a (dependent) type with no additional data, so initiality asserts that any $\Gamma.\mathbf{Void} \vdash A$ type admits a unique displayed algebra homomorphism—a dependent function with no additional conditions—from **Void**.

Likewise, **(Bool, true, false)** is the initial algebra for the signature $X \mapsto \mathbf{1} + \mathbf{1}$. A displayed **(1 + 1)**-algebra over **Bool** is a type $\Gamma.\mathbf{Bool} \vdash A$ type equipped with two terms $\Gamma \vdash a_t : A[\mathbf{id.true}]$ and $\Gamma \vdash a_f : A[\mathbf{id.false}]$; initiality states that for any such displayed algebra there is a unique displayed algebra homomorphism **(Bool, true, false)** $\rightarrow (A, a_t, a_f)$:

$$\rho_{\Gamma, A, a_t, a_f} : \{a \in \mathbf{Tm}(\Gamma.\mathbf{Bool}, A) \mid a_t = a[\mathbf{id.true}] \wedge a_f = a[\mathbf{id.false}]\} \cong \{\star\}$$

We refrain from restating Slogan 2.5.3 in terms of initial algebras, because the general theory of displayed algebras and homomorphisms for a given signature is too significant a detour for this book; we hope that the reader is convinced that a general pattern exists.

Exercise 2.30. In Section 2.5.2, our definition of **Bool** roughly asserted a natural isomorphism between $a \in \text{Tm}(\Gamma.\mathbf{Bool}, A)$ and pairs of substituted terms $(a[\mathbf{id.true}], a[\mathbf{id.false}])$. Prove that this definition is equivalent to the ρ_{Γ,A,a_t,a_f} characterization above.

2.5.4 Unicity via extensional equality

In this section we have defined the inductive types **Void**, **Bool**, and **Nat** by equipping them with constructors and asserting that dependent maps out of them are *judgmentally uniquely determined* by where they send those constructors. That is, a choice of where to send the constructors determines a map via elimination, and any two maps out of an inductive type are judgmentally equal if they agree on the constructors.

This unicity condition is incredibly strong. First of all, it implies the substitution rule for eliminators, because e.g. $\mathbf{if}(a_t, a_f, \mathbf{q})[\gamma.\mathbf{Bool}]$ and $\mathbf{if}(a_t[\gamma], a_f[\gamma], \mathbf{q})$ agree on **true** and **false** (see Exercise 2.26). More alarmingly, in the case of **Void**, it states that *all* terms in contexts containing **Void** are equal to one another (see Exercise 2.25).

It turns out that these unicity principles—the η rules of inductive types—are derivable from the other rules of inductive types in the presence of equality reflection (Section 2.4.4), the other suspiciously strong rule of extensional type theory. For instance:

Theorem 2.5.8. *The following rule (η for **Void**) can be derived from the other rules for **Void** in conjunction with the rules for **Eq**.*

$$\frac{\Gamma \vdash b : \mathbf{Void} \quad \Gamma.\mathbf{Void} \vdash a : A}{\Gamma \vdash \mathbf{absurd}(b) = a[\mathbf{id.b}] : A[\mathbf{id.b}]} \text{Ⓢ,}\Rightarrow$$

Proof. Suppose $\Gamma \vdash b : \mathbf{Void}$ and $\Gamma.\mathbf{Void} \vdash a : A$. By equality reflection (Section 2.4.4), it suffices to exhibit an element of $\mathbf{Eq}(A[\mathbf{id.b}], \mathbf{absurd}(b), a[\mathbf{id.b}])$, which we obtain easily by **Void** elimination:

$$\Gamma \vdash \mathbf{absurd}(b) : \mathbf{Eq}(A, \mathbf{absurd}(b), a[\mathbf{id.b}]) \quad \square$$

In Chapter 3 we will see that all of these suspicious rules are problematic from an implementation perspective, leading us to replace extensional type theory with *intensional type theory* (Chapter 4), which differs formally in only two ways: it replaces **Eq**-types with a different equality type that does not admit equality reflection, and it deletes the η rules from **Void**, **Bool**, and **Nat**.

However, in light of the fact that the latter rules are derivable from the former, we—as is conventional—simply omit the η rules for inductive types from the official specification of extensional type theory. (These rules were all marked as provisional Ⓢ.) Note that this does *not* apply to the η rules for **Π**, **Σ**, **Eq**, or **Unit**, which remain in both type theories.

Semantically, deleting these η rules relaxes the unique existence to simply *existence*. An algebra which admits a (possibly non-unique) algebra homomorphism to any other algebra is known as *weakly initial*, rather than *initial*. Rather than asking for the collection of algebra homomorphisms to be naturally isomorphic to $\{\star\}$, we simply ask for the map from algebra homomorphisms to $\{\star\}$ to admit a natural *section* (right inverse).

Advanced Remark 2.5.9. Recalling Remark 2.5.4, Theorem 2.5.8 corresponds to the fact that a class of morphisms \mathcal{L} which is weakly orthogonal to \mathcal{R} is actually orthogonal to \mathcal{R} when the latter is closed under relative diagonals ($X \rightarrow Y \in \mathcal{R}$ implies $X \rightarrow X \times_Y X \in \mathcal{R}$). \diamond

Exercise 2.31. Prove that the η rule for **Bool** can be derived from the other rules for **Bool** in conjunction with the rules for **Eq**, by mirroring the proof of Theorem 2.5.8.

2.6 Universes: U_0, U_1, U_2, \dots

We are nearly finished with our definition of extensional type theory, but what’s missing is significant: our theory is still not full-spectrum dependent (in the sense of Section 1.1.2)! That is, we have still not introduced the ability to define a family of types whose head type constructor differs at different indices, such as a **Bool**-indexed family of types which sends **true** to **Nat** and **false** to **Unit**. A more subtle but fatal flaw with our current theory is that—despite all the logical connectives at our disposal—we cannot prove that **true** and **false** are different, i.e., we cannot exhibit a term $1 \vdash p : \Pi(\text{Eq}(\text{Bool}, \text{true}, \text{false}), \text{Void})$.

It turns out that addressing the former will solve the latter *en passant*, so in this section we will discuss two approaches for defining dependent types by case analysis. In Section 2.6.1 we introduce *large elimination*, which equips inductive types with a second elimination principle targeting type-valued algebras (which send each constructor to a *type*), in addition to their usual elimination principle targeting algebras valued in a single dependent type (which send each constructor to a *term* of that type).

Unfortunately we will see that large elimination has some serious limitations, so it will not be an official part of our definition of extensional type theory. Instead, in Section 2.6.2, we introduce *type universes*, connectives which internalize the judgment $\Gamma \vdash A$ type modulo “size issues.” By internalizing types as terms of a universe type, universes reduce the problem of computing *types* by case analysis to the problem of computing *terms* by case analysis, which we solved in Section 2.5. That said, universes are a deep and complex topic that will bring us one step closer to our discussion of homotopy type theory in Chapter 5.

2.6.1 Large elimination

In Section 2.5 we introduced elimination principles for inductive types (like **Bool**), which allow us to define dependent functions out of an inductive type ($f : \Pi(\mathbf{Bool}, A)$) by cases on that type’s constructors. A direct but uncommon way of achieving full-spectrum dependency is to equip each inductive type with a second elimination principle, *large elimination*, which allows us to define dependent *families of types* by cases.¹²

In the case of **Bool**, large elimination is characterized by the following rules:

$$\frac{\Gamma \vdash A_t \text{ type} \quad \Gamma \vdash A_f \text{ type} \quad \Gamma \vdash b : \mathbf{Bool}}{\Gamma \vdash \mathbf{If}(A_t, A_f, b) \text{ type}} \text{📎}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A_t \text{ type} \quad \Gamma \vdash A_f \text{ type} \quad \Gamma \vdash b : \mathbf{Bool}}{\Delta \vdash \mathbf{If}(A_t, A_f, b)[\gamma] = \mathbf{If}(A_t[\gamma], A_f[\gamma], b[\gamma]) \text{ type}} \text{📎}$$

$$\frac{\Gamma \vdash A_t \text{ type} \quad \Gamma \vdash A_f \text{ type}}{\Gamma \vdash \mathbf{If}(A_t, A_f, \mathbf{true}) = A_t \text{ type} \quad \Gamma \vdash \mathbf{If}(A_t, A_f, \mathbf{false}) = A_f \text{ type}} \text{📎}$$

If we compare these to the rules of ordinary (“small”) elimination,

$$\frac{\Gamma.\mathbf{Bool} \vdash A \text{ type} \quad \Gamma \vdash a_t : A[\mathbf{id.true}] \quad \Gamma \vdash a_f : A[\mathbf{id.false}] \quad \Gamma \vdash b : \mathbf{Bool}}{\Gamma \vdash \mathbf{if}(a_t, a_f, b) : A[\mathbf{id.b}]}$$

$$\frac{\Gamma.\mathbf{Bool} \vdash A \text{ type} \quad \Gamma \vdash a_t : A[\mathbf{id.true}] \quad \Gamma \vdash a_f : A[\mathbf{id.false}]}{\Gamma \vdash \mathbf{if}(a_t, a_f, \mathbf{true}) = a_t : A[\mathbf{id.true}] \quad \Gamma \vdash \mathbf{if}(a_t, a_f, \mathbf{false}) = a_f : A[\mathbf{id.false}]}$$

we see that the large eliminator **If** is exactly analogous to the small eliminator **if** “specialized to $A := \text{type}$.” Note that this statement is nonsense because the judgment “type” is not a type, but the intuition is useful and will be formalized momentarily. (Indeed, for this reason we cannot formally obtain **If** as a special case of **if**.) Continuing on with the metaphor, the rule for **If** is simpler than the rule for **if** because it has a fixed codomain “type” which is moreover *not* dependent on **Bool**: it makes no sense to ask for “ $\Gamma \vdash A_t \text{ type}[\mathbf{id.true}]$.”

It is even more standard to omit the η rule for large elimination than for small elimination (which is itself typically omitted), but such a rule would state that dependent types indexed by **Bool** are uniquely determined by their values on **true** and **false**:

$$\frac{\Gamma.\mathbf{Bool} \vdash A \text{ type} \quad \Gamma \vdash b : \mathbf{Bool}}{\Gamma \vdash A[\mathbf{id.b}] = \mathbf{If}(A[\mathbf{id.true}], A[\mathbf{id.false}], b) \text{ type}} \text{📎📎}$$

¹²Large elimination maps **Bool** into the collection of all types, which is “large” (in the sense of being “the proper class of all sets”) rather than the collection of terms of a single type, which is “small” (“a set”).

If we include the η rule, then the rules for **If** would express that instantiating a **Bool**-indexed type at **true** and **false**, i.e. $((\text{id.true})^*, (\text{id.false})^*)$, has a natural inverse:

$$((\text{id.true})^*, (\text{id.false})^*) : \text{Ty}(\Gamma.\mathbf{Bool}) \cong \text{Ty}(\Gamma) \times \text{Ty}(\Gamma)$$

Again, compare this to our original formulation of small elimination for **Bool**:

$$((\text{id.true})^*, (\text{id.false})^*) : \text{Tm}(\Gamma.\mathbf{Bool}, A) \cong \text{Tm}(\Gamma, A[\text{id.true}]) \times \text{Tm}(\Gamma, A[\text{id.false}])$$

When we elide η , large elimination instead states that this map has a *section* (a right inverse), which is to say that a choice of where to send **true** and **false** determines a family of types via **If**, but *not uniquely*. This follows the discussion in Section 2.5.4, except that we *cannot* derive the η rule for large elimination from extensional equality because there is no type “ $\text{Eq}(\text{type}, -, -)$ ” available to carry out the argument in Theorem 2.5.8.

Remark 2.6.1. Large elimination only applies to types defined by mapping-out properties such as inductive types; there is no corresponding principle for mapping-in connectives like $\Pi(A, B)$ because these do not quantify over any target, whether “small” or “large.” \diamond

Remark 2.6.2. If we have both small and large elimination for **Bool**, then we can combine them into a derived induction principle for **Bool** that works for any $a_t : A_t$ and $a_f : A_f$, using large elimination to define the type family into which we perform a small elimination.

$$\frac{\Gamma \vdash a_t : A_t \quad \Gamma \vdash a_f : A_f \quad \Gamma \vdash b : \mathbf{Bool}}{\Gamma \vdash \text{if}(a_t, a_f, b) : \mathbf{If}(A_t, A_f, b)} \quad \text{e}_\Rightarrow \quad \diamond$$

With large elimination—or a related feature, type universes—we can prove the disjointness of the booleans. (Although the proof below uses equality reflection, the same theorem holds in intensional type theory for essentially the same reason.) Our claim that we *cannot* prove disjointness without these features is a (relatively simple) independence metatheorem requiring a model construction; see *The Independence of Peano’s Fourth Axiom from Martin-Löf’s Type Theory Without Universes* [Smi88].

Theorem 2.6.3. *Using the rules for **If**, there is a term*

$$1 \vdash \text{disjoint} : \Pi(\text{Eq}(\mathbf{Bool}, \text{true}, \text{false}), \mathbf{Void})$$

Proof. We informally describe the derivation of **disjoint**. By Π -introduction we may assume $\text{Eq}(\mathbf{Bool}, \text{true}, \text{false})$ and prove **Void**. In order to do this, consider the following auxiliary family of types over **Bool**:

$$1.\text{Eq}(\mathbf{Bool}, \text{true}, \text{false}).\mathbf{Bool} \vdash P := \text{If}(\mathbf{Unit}, \mathbf{Void}, q) \text{ type}$$

Then

$$\begin{aligned}
 1.\text{Eq}(\text{Bool}, \text{true}, \text{false}) &\vdash \text{Unit} \\
 &= P[\text{id.true}] && \text{by } \beta \text{ for If} \\
 &= P[\text{id.false}] && \text{by equality reflection on } \mathbf{q} \\
 &= \text{Void type} && \text{by } \beta \text{ for If}
 \end{aligned}$$

and therefore $1.\text{Eq}(\text{Bool}, \text{true}, \text{false}) \vdash \text{tt} : \text{Void}$. In sum, we define $\text{disjoint} := \lambda(\text{tt})$. \square

As for other inductive types, the large elimination principle of **Void** is:

$$\frac{\Gamma \vdash a : \text{Void}}{\Gamma \vdash \text{Absurd}(a) \text{ type}} \text{ }$$

Unfortunately, we run into a problem when trying to define large elimination for **Nat**.

$$\frac{\Gamma \vdash n : \text{Nat} \quad \Gamma \vdash A_z \text{ type} \quad \Gamma.\text{Nat}.\text{“type”} \vdash A_s \text{ type}}{\Gamma \vdash \text{Rec}(A_z, A_s, n) \text{ type}} \text{ } !?$$

In the ordinary eliminator, the branch for $\text{succ}(-)$ has two variables $m : \text{Nat}$, $a : A(m)$ binding the predecessor m and (recursively) the result a of the eliminator on m . When “ $A := \text{type}$ ” the recursive result is a *type*, meaning that the $\text{succ}(-)$ branch ought to bind a *type variable*, a concept which is not a part of our theory. This is a serious problem because recursive constructions of types were a major class of examples in Section 1.1.2.

Exercise 2.32. There is however a *non-recursive* large elimination principle for **Nat** which defines a type by case analysis on whether a number is **zero**. This principle follows from the rules in this section along with the other rules of extensional type theory; state and define it.

2.6.2 Universes

Although large elimination is a useful concept, it sees essentially no use in practice. We have just seen one reason: large eliminators cannot be recursive. The standard approach is instead to include *universe types*, which are “types of types,” or types which internalize the judgment $\Gamma \vdash A \text{ type}$. Using universes, we can recover large elimination as small elimination into a universe; we are also able to express polymorphic type quantification using dependent functions out of a universe.

A universe is a type with no parameters, so its formation rule is once again a natural family of constants $U_\Gamma \in \text{Ty}(\Gamma)$, or

$$\frac{}{\Gamma \vdash \mathbf{U} \text{ type}} \quad \frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{U} = \mathbf{U}[\gamma] \text{ type}}$$

As for its terms, the most straightforward definition would be to stipulate a natural isomorphism between terms of U and types:

$$\iota : \text{Tm}(\Gamma, U) \cong \text{Ty}(\Gamma) \quad (?!)$$

Note that just as we did not ask for terms of Π -types to literally be terms with an extra free variable, we cannot ask for terms of U to literally be types: these are two different sorts!

In inference rules, the forward map of the isomorphism would introduce a new type former $\text{El}(-)$ ¹³ which “decodes” an element of U into a genuine type. The reverse map conversely “encodes” a genuine type as an element of U . These intuitions lead us to often refer to elements of U as *codes* for types.

$$\frac{\Gamma \vdash a : U}{\Gamma \vdash \text{El}(a) \text{ type}} \qquad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : U}{\Delta \vdash \text{El}(a)[\gamma] = \text{El}(a[\gamma]) \text{ type}}$$

$$\frac{\Gamma \vdash A \text{ type}}{\Gamma \vdash \text{code}(A) : U} \text{?!} \qquad \frac{\Gamma \vdash A \text{ type}}{\Gamma \vdash \text{El}(\text{code}(A)) = A \text{ type}} \text{?!} \qquad \frac{\Gamma \vdash c : U}{\Gamma \vdash \text{code}(\text{El}(c)) = c : U} \text{?!}$$

Unfortunately we can’t have nice things, as the last three rules above—the ones involving **code**—are unsound. In particular they imply that U contains (a code for) U , making it a “type of all types, including itself” and therefore subject to a variation on Russell’s paradox known as *Girard’s paradox* [Coq86], as outlined in Section 2.7.

2.6.2.1 Populating the universe

Returning to our definition of universe types, it is safe to postulate a type U of type-codes which decode via El into types. (Indeed, with large elimination it is even possible to define such a type manually, e.g. $U := \text{Bool}$ with $\text{El}(\text{true}) := \text{Unit}$ and $\text{El}(\text{false}) := \text{Void}$.)

$$U_\Gamma \in \text{Ty}(\Gamma)$$

$$\text{El} : \text{Tm}(\Gamma, U) \rightarrow \text{Ty}(\Gamma)$$

Our first attempt at populating $\text{Tm}(\Gamma, U)$ was to ask for an inverse to El , but that turns out to be inconsistent. Instead, we will simply manually equip U with codes decoding to the type formers we have presented so far, but crucially *not* with a code for U itself. This approach is somewhat verbose—for each type former we add an introduction rule for U , a substitution rule, and an equation stating that El decodes it to the corresponding type—but it allows us to avoid Girard’s paradox while still populating U with codes for (almost) every type in our theory.

¹³This name is not so mysterious: it means “*elements* of,” and is pronounced “ell” or, often, omitted.

For example, to close U under dependent function types we add the following rules:

$$\frac{\Gamma \vdash a : U \quad \Gamma.\text{El}(a) \vdash b : U}{\Gamma \vdash \text{pi}(a, b) : U} \quad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : U \quad \Gamma.\text{El}(a) \vdash b : U}{\Delta \vdash \text{pi}(a, b)[\gamma] = \text{pi}(a[\gamma], b[\gamma.\text{El}(a)]) : U}$$

$$\frac{\Gamma \vdash a : U \quad \Gamma.\text{El}(a) \vdash b : U}{\Gamma \vdash \text{El}(\text{pi}(a, b)) = \Pi(\text{El}(a), \text{El}(b)) \text{ type}}$$

The third rule states that $\text{pi}(a, b)$ is the code in U for the type $\Pi(\text{El}(a), \text{El}(b))$. Note that the context of b in the introduction rule for $\text{pi}(a, b)$ makes reference to $\text{El}(a)$, mirroring the dependency structure of Π -types. Although this move is forced, it means that the definitions of U and El each reference the other—the type of a constructor of U mentions El , and the type of El itself mentions U —so U and El must be defined simultaneously. In fact, this is the paradigmatic example of an *inductive-recursive* definition, an inductive type that is defined simultaneously with a recursive function out of it [Dyb00].

It is no more difficult to close U under dependent pairs, extensional equality, the unit type, and inductive types. These rules quickly become tedious, so we write only their introduction rules below, leaving the remaining rules to Appendix A.

$$\frac{\Gamma \vdash a : U \quad \Gamma.\text{El}(a) \vdash b : U}{\Gamma \vdash \text{sig}(a, b) : U} \quad \frac{\Gamma \vdash a : U \quad \Gamma \vdash x, y : \text{El}(a)}{\Gamma \vdash \text{eq}(a, x, y) : U}$$

$$\frac{}{\Gamma \vdash \text{unit} : U} \quad \frac{}{\Gamma \vdash \text{nat} : U} \quad \frac{}{\Gamma \vdash \text{void} : U} \quad \frac{}{\Gamma \vdash \text{bool} : U}$$

We can now recover the large elimination principles of Section 2.6.1 in terms of small elimination into the type U . Moreover, because we can perfectly well extend the context by a variable of type U , we can now also construct types by recursion on natural numbers:

$$\frac{\Gamma \vdash n : \text{Nat} \quad \Gamma \vdash a_z : U \quad \Gamma.\text{Nat}.U \vdash a_s : U}{\Gamma \vdash \text{Rec}(a_z, a_s, n) := \text{El}(\text{rec}(a_z, a_s, n)) \text{ type}} \Rightarrow$$

Notation 2.6.4. In general, we may refer to a pair of a type ($\Gamma \vdash B$ type) and a type family over that type ($\Gamma.B \vdash E$ type) as a *universe*, if it is appropriate to think of the former as a collection of codes and the latter a decoding function, generalizing the paradigmatic example of (U, El) . We will encounter other universes in Sections 2.8 and 5.2.

Remark 2.6.5. Proof assistant users are very familiar with universes, so such readers may be wondering why they have never seen El before. Indeed, proof assistants such as Coq and Agda treat types and elements of U as indistinguishable. Historically, much of the literature calls such universes—for which $\text{Tm}(\Gamma, U) \subseteq \text{Ty}(\Gamma)$ —*universes à la Russell*, in contrast to our *universes à la Tarski*, but we find such a subset inclusion to be meta-suspicious.

Instead, we prefer to say that Coq and Agda programs do not expose the notion of type to the user at all, instead consistently referring only to elements of U . This obviates the need for the user to ever write or see \mathbf{El} , and the necessary calls to \mathbf{El} can be inserted automatically by the proof assistant in a process known as *elaboration*. \diamond

Remark 2.6.6. Another more semantically natural variation of universes relaxes the judgmental equalities governing \mathbf{El} to *isomorphisms* $\mathbf{El}(\mathbf{pi}(a, b)) \cong \Pi(\mathbf{El}(a), \mathbf{El}(b))$, producing what are known as *weak universes à la Tarski*. However, our *strict* formulation is more standard and more convenient. \diamond

Advanced Remark 2.6.7. Universes in type theory play a similar role to Grothendieck universes and their categorical counterparts in set theory and category theory. We often refer to types encoded by U as *small* or U -small, and ask for small types to be closed under various operations. As a result, universes in type theory roughly have the same proof-theoretical strength as strongly inaccessible cardinals. Note, however, that the lack of choice and excluded middle in type theory precludes a naïve comparison between it and ZFC or similar theories; see Section 3.5.1. \diamond

2.6.3 Hierarchies of universes

Our definition of U is perfectly correct, but the fact that U lacks a code for itself means that we cannot recursively define types that mention U . In addition, although we can quantify over “small” types with $\Pi(U, -)$, we cannot write any type quantifiers whose domain includes U . We cannot fix these shortcomings directly, but we can mitigate them by defining a *second* universe type U_1 closed under all the same type codes as before *as well as a code for* U , but no code for U_1 itself. The same problem occurs one level up, so we add a third universe U_2 containing codes for U and U_1 but not U_2 , and so forth.

In practice, nearly all applications of type theory require only a finite number of universes, but for uniformity and because this number varies between applications, it is typical to ask for a countably infinite (alternatively, finite but arbitrary) tower of universes each of which contains codes for the smaller ones. (For uniformity we write $U_0 := U$.) This collection of U_i is known as a *universe hierarchy*.

To define an infinite number of types and terms, we must now write *rule schemas*, collections of rules that must be repeated for every (external, not internal) natural number $i > 1$. Each of these rules follows the same pattern in U , with one new feature: U_i contains

a code $\mathbf{uni}_{i,j}$ for \mathbf{U}_j whenever j is strictly smaller than i .

$$\begin{array}{c}
\frac{}{\Gamma \vdash \mathbf{U}_i \text{ type}} \qquad \frac{\Gamma \vdash a : \mathbf{U}_i}{\Gamma \vdash \mathbf{El}_i(a) \text{ type}} \qquad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : \mathbf{U}_i}{\Delta \vdash \mathbf{El}_i(a)[\gamma] = \mathbf{El}_i(a[\gamma]) \text{ type}} \\
\\
\frac{\Gamma \vdash a : \mathbf{U}_i \quad \Gamma.\mathbf{El}_i(a) \vdash b : \mathbf{U}_i}{\Gamma \vdash \mathbf{pi}_i(a, b) : \mathbf{U}_i} \quad \frac{\Gamma \vdash a : \mathbf{U}_i \quad \Gamma \vdash x, y : \mathbf{El}_i(a)}{\Gamma \vdash \mathbf{eq}_i(a, x, y) : \mathbf{U}_i} \\
\\
\frac{}{\Gamma \vdash \mathbf{unit}_i : \mathbf{U}_i} \quad \frac{}{\Gamma \vdash \mathbf{void}_i : \mathbf{U}_i} \quad \frac{}{\Gamma \vdash \mathbf{bool}_i : \mathbf{U}_i} \quad \frac{}{\Gamma \vdash \mathbf{nat}_i : \mathbf{U}_i} \quad \frac{j < i}{\Gamma \vdash \mathbf{uni}_{i,j} : \mathbf{U}_i}
\end{array}$$

Again for uniformity we write $\mathbf{pi}_0(a, b) := \mathbf{pi}(a, b)$, etc., and we omit the substitution rules for type codes as well as the type equations explaining how each \mathbf{El}_i computes on codes, such as $\mathbf{El}_i(\mathbf{eq}_i(a, x, y)) = \mathbf{Eq}(\mathbf{El}_i(a), x, y)$ and $\mathbf{El}_i(\mathbf{uni}_{i,j}) = \mathbf{U}_j$.

It is easy to see that the rules for \mathbf{U}_{i+1} are a superset of the rules for \mathbf{U}_i : the only difference is the addition of the code $\mathbf{uni}_{i+1,i} : \mathbf{U}_{i+1}$ and codes that mention this code, such as $\mathbf{pi}_{i+1}(\mathbf{uni}_{i+1,i}, \mathbf{uni}_{i+1,i}) : \mathbf{U}_{i+1}$. Thus it should be possible to prove that every closed code of type \mathbf{U}_i has a counterpart of type \mathbf{U}_{i+1} that decodes to the same type, that is, “ $\mathbf{U}_i \subseteq \mathbf{U}_{i+1}$.” However, this fact is not visible inside the theory. We have no induction principle for the universe, so we cannot define an “inclusion” function $f : \mathbf{U}_i \rightarrow \mathbf{U}_{i+1}$ much less prove that it satisfies $\mathbf{El}_{i+1}(f(a)) = \mathbf{El}_i(a)$. And there is simply no way, external or otherwise, to “lift” a variable of type \mathbf{U}_i to the type \mathbf{U}_{i+1} .

We thus equip our universe hierarchy with one final operation: a *lifting* operation that includes elements of \mathbf{U}_i into \mathbf{U}_{i+1} , which is compatible with \mathbf{El} and sends type codes of \mathbf{U}_i to their counterparts in \mathbf{U}_{i+1} . Such a strict lifting operation allows users to generally avoid worrying about universe levels, because small codes can always be hoisted up to their larger counterparts when needed.

$$\begin{array}{c}
\frac{\Gamma \vdash c : \mathbf{U}_i}{\Gamma \vdash \mathbf{lift}_i(c) : \mathbf{U}_{i+1}} \qquad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : \mathbf{U}_i}{\Delta \vdash \mathbf{lift}_i(a)[\gamma] = \mathbf{lift}_i(a[\gamma]) : \mathbf{U}_{i+1}} \\
\\
\frac{\Gamma \vdash a : \mathbf{U}_i}{\Gamma \vdash \mathbf{El}_{i+1}(\mathbf{lift}_i(a)) = \mathbf{El}_i(a) \text{ type}}
\end{array}$$

The last rule above states that a code and its lift both encode the same type. Recalling that the entire point of a universe hierarchy is to get as close as possible to “ $\mathbf{U} : \mathbf{U}$ ” without being inconsistent, it makes sense to treat lifts as a clerical operation that does not affect the type about which we speak. In addition, this equation is actually needed to state that

lift commutes with codes, such as **pi** (other rules omitted):

$$\frac{\Gamma \vdash a : U_i \quad \Gamma.\mathbf{El}_i(a) \vdash b : U_i}{\Gamma \vdash \mathbf{lift}_i(\mathbf{pi}_i(a, b)) = \mathbf{pi}_{i+1}(\mathbf{lift}_i(a), \mathbf{lift}_i(b)) : U_{i+1}}$$

Remark 2.6.8. We say a universe hierarchy is (strictly) *cumulative* when it is equipped with **lift** operations that commute (strictly) with codes. Historically the term “cumulativity” often refers to material subset inclusions $\text{Tm}(\Gamma, U_i) \subseteq \text{Tm}(\Gamma, U_{i+1})$ but once again such conditions are incompatible with our perspective. \diamond

Remark 2.6.9. There is an equivalent presentation of universe hierarchies known as *universes à la Coquand* in which one stratifies the type judgment itself, and the i th universe precisely internalizes the i th type judgment [Coq13; Coq19; Gra+21; FAM23]. That is, we have sorts $\text{Ty}_i(\Gamma)$ for $i \in \mathbb{N} + \{\top\}$ with $\text{Ty}(\Gamma) := \text{Ty}_\top(\Gamma)$, and natural isomorphisms $\text{Ty}_i(\Gamma) \cong \text{Tm}(\Gamma, U_i)$ for $i \in \mathbb{N}$ mediated by $\mathbf{El}_i/\mathbf{code}_i$. This presentation essentially creates a new judgmental structure designed to be internalized by \mathbf{U} , and has the concrete benefit of unifying type formation and universe introduction into a single set of rules. \diamond

Exercise 2.33. Check that the equational rule $\mathbf{lift}_i(\mathbf{pi}_i(a, b)) = \mathbf{pi}_{i+1}(\mathbf{lift}_i(a), \mathbf{lift}_i(b))$ above is meta-well-typed. (Hint: you need to use $\mathbf{El}_{i+1}(\mathbf{lift}_i(a)) = \mathbf{El}_i(a)$.)

Exercise 2.34. We only included lifts from U_i to U_{i+1} , rather than from U_i to U_j for every $i < j$. Show that the latter notion of lift is derivable for any concrete $i < j$ and that it satisfies the expected equations.

2.7★ Girard's paradox

We now substantiate the claim in Section 2.6.2 that it is inconsistent for \mathbf{U} to contain a code for itself, a fact commonly known as Girard's paradox; specifically, we present a simplified argument due to Hurkens [Hur95].¹⁴ The details of this paradox are not relevant to any later material in this book, so the reader may freely skip this section. In this section alone, we adopt the (inconsistent) rules of Section 2.6.2 pertaining to **code**.

At a high level, the fact that \mathbf{U} contains a code for itself means that we can construct a universe Θ that admits an embedding from its own double power set $\mathcal{P}(\mathcal{P}\Theta)$; from this we can define a “set of all ordinals” and carry out a version of the Burali-Forti paradox. The details become somewhat involved, in part because the standard paradoxes of set theory rely on comprehension and extensionality principles not available to us in type

¹⁴An Agda formalization of Hurkens's paradox is available at <https://github.com/agda/agda/blob/master/test/Succeed/Hurkens.agda>; formalizations in other proof assistants are readily available online.

theory. Indeed, historically it was far from clear that “ $U : U$ ” was inconsistent, and in fact Martin-Löf's first version of type theory had this very flaw [ML71].

$$\begin{aligned} \mathcal{P} &: U \rightarrow U \\ \mathcal{P} A &= \mathbf{code}(\mathbf{El}(A) \rightarrow U) \end{aligned}$$

$$\begin{aligned} \mathcal{P}^2 &: U \rightarrow U \\ \mathcal{P}^2 A &= \mathcal{P} (\mathcal{P} A) \end{aligned}$$

$$\begin{aligned} \Theta &: U \\ \Theta &= \mathbf{code}((A : U) \rightarrow (\mathbf{El}(\mathcal{P}^2 A) \rightarrow \mathbf{El}(A)) \rightarrow \mathbf{El}(\mathcal{P}^2 A)) \end{aligned}$$

Lemma 2.7.1 (Powerful universe). *The universe Θ admits maps*

$$\begin{aligned} \tau &: \mathbf{El}(\mathcal{P}^2 \Theta) \rightarrow \Theta \\ \sigma &: \Theta \rightarrow \mathbf{El}(\mathcal{P}^2 \Theta) \end{aligned}$$

such that

$$(C : \mathbf{El}(\mathcal{P}^2 \Theta)) \rightarrow (\sigma (\tau C) = \lambda(\phi : \mathbf{El}(\mathcal{P} \Theta)) \rightarrow C(\phi \circ \tau \circ \sigma))$$

Proof. We define:

$$\begin{aligned} \tau &: \mathbf{El}(\mathcal{P}^2 \Theta) \rightarrow \mathbf{El}(\Theta) \\ \tau (\Phi : \mathbf{El}(\mathcal{P}^2 \Theta)) (A : U) (f : \mathbf{El}(\mathcal{P}^2 A) \rightarrow \mathbf{El}(A)) (\chi : \mathbf{El}(\mathcal{P} A)) &= \\ &\Phi (\lambda(\theta : \Theta) \rightarrow \chi (f (\theta A f))) \end{aligned}$$

$$\begin{aligned} \sigma &: \mathbf{El}(\Theta) \rightarrow \mathbf{El}(\mathcal{P}^2 \Theta) \\ \sigma \theta &= \theta \Theta \tau \end{aligned}$$

We leave the equational condition to Exercise 2.35. □

Exercise 2.35. Show that the above definitions of τ and σ satisfy the necessary equation.

As an immediate consequence of Lemma 2.7.1, we have:

$$\sigma (\tau (\sigma x)) = \lambda(\phi : \mathbf{El}(\mathcal{P} \Theta)). \sigma x (\phi \circ \tau \circ \sigma) \tag{2.1}$$

One way to understand the statement of Lemma 2.7.1 is that, regarding \mathcal{P} as a functor whose action on $f : \mathbf{El}(Y) \rightarrow \mathbf{El}(X)$ is precomposition $f^* : \mathbf{El}(\mathcal{P} X) \rightarrow \mathbf{El}(\mathcal{P} Y)$, the equational condition is equivalent to $\sigma \circ \tau = (\tau \circ \sigma)^{**}$.

We derive a contradiction from Lemma 2.7.1 by constructing ordinals within Θ :

-- $y < x$ (“ $y \in x$ ”) when each f in σx contains y
 $(<) : \mathbf{El}(\Theta) \rightarrow \mathbf{El}(\Theta) \rightarrow \mathbf{U}$
 $y < x = \mathbf{code}((f : \mathbf{El}(\mathcal{P} \Theta)) \rightarrow \mathbf{El}(\sigma x f) \rightarrow \mathbf{El}(f y))$

-- f is inductive if for all x , if f is in σx then x is in f
 $\mathbf{ind} : \mathbf{El}(\mathcal{P} \Theta) \rightarrow \mathbf{U}$
 $\mathbf{ind} f = \mathbf{code}((x : \mathbf{El}(\Theta)) \rightarrow \mathbf{El}(\sigma x f) \rightarrow \mathbf{El}(f x))$

-- x is well-founded if it is in every inductive f
 $\mathbf{wf} : \mathbf{El}(\Theta) \rightarrow \mathbf{U}$
 $\mathbf{wf} x = \mathbf{code}((f : \mathbf{El}(\mathcal{P} \Theta)) \rightarrow \mathbf{El}(\mathbf{ind} f) \rightarrow \mathbf{El}(f x))$

Specifically, we consider $\Omega := \tau(\lambda f \rightarrow \mathbf{ind} f)$, the collection of all inductive collections. Using Lemma 2.7.1 we argue that Ω is both well-founded and not well-founded.

Lemma 2.7.2. Ω is well-founded.

Proof. Suppose $f : \mathbf{El}(\mathcal{P} \Theta)$ is inductive; we must show $\mathbf{El}(f \Omega)$. By the definition of \mathbf{ind} , for this it suffices to show $\mathbf{El}(\sigma \Omega f)$. Unfolding the definition of Ω and rewriting by the equation in Lemma 2.7.1 with $C := \mathbf{ind}$, it suffices to show that $f \circ \tau \circ \sigma$ is inductive.

Thus suppose we are given $x : \mathbf{El}(\Theta)$ such that $\mathbf{El}(\sigma x (f \circ \tau \circ \sigma))$; we must show $\mathbf{El}(f (\tau (\sigma x)))$. By rewriting $\mathbf{El}(\sigma x (f \circ \tau \circ \sigma))$ along Equation (2.1), we conclude that $\mathbf{El}(\sigma (\tau (\sigma x)) f)$. However, by our assumption that f is inductive, this implies $\mathbf{El}(f (\tau (\sigma x)))$, which is what we wanted to show. \square

To prove that Ω is not well-founded, we start by showing that the collection of “sets not containing themselves” $\phi := \lambda y \rightarrow \mathbf{code}(\mathbf{El}(\tau (\sigma y) < y) \rightarrow \mathbf{Void})$ is inductive.

Lemma 2.7.3. ϕ is inductive.

Proof. Suppose we are given x such that $\mathbf{El}(\sigma x \phi)$; we must show $\mathbf{El}(\tau (\sigma x) < x) \rightarrow \mathbf{Void}$. Thus suppose $\mathbf{El}(\tau (\sigma x) < x)$, which is to say that for any f such that $\mathbf{El}(\sigma x f)$, we have $\mathbf{El}(f (\tau (\sigma x)))$. Using our hypothesis we may set $f := \phi$, from which we conclude $\mathbf{El}(\tau (\sigma (\tau (\sigma x))) < \tau (\sigma x)) \rightarrow \mathbf{Void}$. We derive the required contradiction by proving that $\mathbf{El}(\tau (\sigma (\tau (\sigma x))) < \tau (\sigma x))$ holds, by $\mathbf{El}(\tau (\sigma x) < x)$ and Exercise 2.36. \square

Exercise 2.36. Show that $\text{El}(x < y)$ implies $\text{El}(\tau(\sigma x) < \tau(\sigma y))$.

Theorem 2.7.4. *Void is inhabited.*

Proof. Because Ω is well-founded and ϕ is inductive, we have $\text{El}(\tau(\sigma \Omega) < \Omega) \rightarrow \mathbf{Void}$. To derive a contradiction, it suffices to show $\text{El}(\tau(\sigma \Omega) < \Omega)$, which is to say that for any f such that $\text{El}(\sigma \Omega f)$, we have $\text{El}(f(\tau(\sigma \Omega)))$. By the definition of Ω , $\text{El}(\sigma(\Omega f))$ implies that $f \circ \tau \circ \sigma$ is inductive; combining this with the fact that Ω is well-founded, we obtain $\text{El}(f(\tau(\sigma \Omega)))$ as required. \square

2.8[★] Propositions and universes of propositions (DRAFT)

Many readers will have encountered the slogan “propositions are types”, perhaps termed the Curry–Howard correspondence or the Brouwer–Heyting–Kolmogorov interpretation. This dictum states that that various connectives of type theory should be possible to simultaneously view as program specifications and logical operations. For dependent type theory, this perspective was strongly espoused by Martin-Löf [ML82].

We will not fully develop this theme here for reasons of space, but to motivate the forthcoming discussion it is helpful to develop a few examples. First, Π -types in dependent type theory link together functions (à la functional programming) and universal quantification/implication from logic. While there are many ways of making this linkage precise, the most basic is to simply observe that the typing rules for Π -types in the non-dependent case are identical to those for implication in intuitionistic propositional logic, to wit:

$$\frac{\Gamma.A \vdash b : B[\mathbf{p}]}{\Gamma \vdash \lambda(b) : A \rightarrow B} \quad \text{versus} \quad \frac{\Gamma, p \vdash q \text{ true}}{\Gamma \vdash p \rightarrow q \text{ true}}$$

$$\frac{\Gamma \vdash f : A \rightarrow B \quad \Gamma \vdash a : A}{\Gamma \vdash \mathbf{app}(f, a) : B} \quad \text{versus} \quad \frac{\Gamma \vdash p \rightarrow q \text{ true} \quad \Gamma \vdash p \text{ true}}{\Gamma \vdash q \text{ true}}$$

The same observations can be made for (non-dependent) Σ -types and conjunction, for \mathbf{Unit} and \top , and for \mathbf{Void} and \perp . If we allow for dependence, then connectives like Π - and Σ -types resemble *quantifiers* rather than simple connectives and equality types capture the equality predicate of first-order logic.

This correspondence, however, is already slightly imperfect: in logic, one uses \forall to quantify over elements of a sort (e.g., to quantify over sets when axiomatizing set theory) and implication to denote hypothesizing over the truth of a proposition. These have collapsed into a single operation in dependent type theory. Equivalently, one might say that type theory does not distinguish between sorts and propositions by regarding propositions as the sort of their proofs.

Thus, while one can embed propositional logic into type theory by reading $\Gamma \vdash a : A$ as “ a is a proof of the proposition A ” rather than “ a is a program of type A ”, this is only half the picture when upgrading the correspondence to first-order logic. There one must also sometimes read $\Gamma \vdash a : A$ as “ a is an element of the sort A ”. One can see both of these interpretations on display when we translate $\forall x : \tau. \phi(x)$ to $\Pi(A, B)$ as A is used to interpret a sort and B a proposition. We refer the reader to Martin-Löf [ML84b] or Girard, Taylor, and Lafont [GTL89] for thorough descriptions of this correspondence.

Many types have an intrinsic bias as to whether they are best viewed as “sorts” or “propositions”. For example, if we regard **Nat** or **Bool** as propositions, we find that they are interchangeable: it is possible to inhabit both these types and, generally, logic is not concerned with *how* a particular proposition is proven. Of course, the rules governing **Nat** and **Bool** are quite different and so regarding them as both merely “true” misses out on much of their behavior. On the other hand, we have encountered types—**Unit** and **Void**—where the elements of these types were completely interchangeable and are therefore naturally viewed as propositions.

In this section we explore when a type is best regarded as a proposition along with the special role played by such types. We in particular discuss the addition of a universe Ω whose elements are codes for propositions and the consequences of this addition. Finally, we discuss the character of “logic within type theory” and observe that many classical tautologies from logic do not hold within type theory.

Warning 2.8.1. While we discuss several new constructions and types in this section, none of them will be considered part of our official definition of extensional type theory. Moreover, this material is not used in the remainder of the book with the exception of a passing mention in Section 5.1.

2.8.1 Propositions in type theory

We begin by observing that the distinction between proposition and a “proper” type is quite important. Consider, for instance, the axiom of choice:

$$(\forall a : A. \exists b : B. P(a, b)) \Rightarrow (\exists f : A \rightarrow B. \forall a : A. P(a, f(a)))$$

One can naively translate this into dependent type theory by replacing \forall with Π , \exists with Σ , and the predicate P with a type family $A \times B \rightarrow \mathbf{U}$. Under this translation, however, this “type-theoretic axiom of choice” is trivial!

Exercise 2.37. Construct an element of the following type:

TTChoice =

$$((a : A) \rightarrow \sum_{b:B} \mathbf{El}(P(a, b))) \rightarrow (\sum_{f:A \rightarrow B} (a : A) \rightarrow \mathbf{El}(P(a, f(a))))$$

In light of this, two conclusions are possible: either (1) type theory is so powerful as to settle the debate on the status of the axiom of choice or (2) this type-theoretic axiom of choice fails to capture some important aspects of the logical formulation. Unfortunately, the latter conclusion is more persuasive.

Indeed, there is a vital difference between how one ought to interpret the proposition $\exists A. B$ and the type $\Sigma(A, B)$: with Σ we are given a particular choice of the element of A used to validate B while with \exists we have no such selection. In particular, if there are likely to be many, many distinct elements of $\Sigma(A, B)$ and our subsequent constructions can behave differently for each such element. On the other hand, there can only ever be one proof of $\exists A. B$ and its construction cannot influence subsequent arguments. This is the reason why the type-theoretic axiom of choice trivializes, while the set-theoretic version is quite rich and complex.

For this reason, we might refine our slogan from “propositions are types” to “propositions are types... but not all types are propositions”. This leads us to the question of which types, exactly, can be regarded as propositions and to the following definition:

Definition 2.8.2. $\Gamma \vdash A$ type is said to be a proposition when the following holds:

$$\Gamma.A.A[\mathbf{p}] \vdash \mathbf{q}[\mathbf{p}] = \mathbf{q} : A[\mathbf{p}^2]$$

More informally, A is a proposition if all elements of A are equal.

We can reformulate this requirement using `Eq`: a type $\Gamma \vdash A$ type is a proposition just when `isPropA = $\Gamma \vdash \Pi(A, \Pi(A[\mathbf{p}], \text{Eq}(A[\mathbf{p}^2], \mathbf{q}[\mathbf{p}], \mathbf{q})))$` type is inhabited.

Exercise 2.38. Prove that `isPropA` coincides with the given definition of propositions.

Exercise 2.39. Prove that $\Gamma \vdash A$ type is a proposition if and only if the following substitution is an isomorphism: $\Gamma.A.A[\mathbf{p}] \vdash \mathbf{p} : \Gamma.A$.

Exercise 2.40. Prove $\Gamma \vdash A$ type is a proposition iff for all $\Delta \vdash \gamma : \Gamma$, $|\text{Tm}(\Delta, A[\gamma])| \leq 1$.

We do not have to look far to find both examples and non-examples of propositions within type theory. For instance, `Unit` and `Void` are both propositions while `Bool` and `Nat` are not (Theorem 2.6.3).

Exercise 2.41. Given $\Gamma \vdash a, b : A$ show that $\Gamma \vdash \text{Eq}(A, a, b)$ type is a proposition.

Exercise 2.42. Given $\Gamma \vdash A$ type and $\Gamma.A \vdash B$ type such that the latter is a proposition, show that $\Gamma \vdash \Pi(A, B)$ type is a proposition.

Exercise 2.43. Show that if $\Gamma \vdash A$ type and $\Gamma.A \vdash B$ type are both propositions, so too is $\Gamma \vdash \Sigma(A, B)$ type.

This last point sheds further light on the distinction between the type-theoretic axiom of choice and the logical axiom of choice: if A is not a proposition then $\Sigma(A, B)$ behaves quite differently than $\exists A. B$, *even if B is a proposition*. We now set out to give a more faithful encoding of the axiom of choice within type theory and, eventually, discuss its validity within type theory.

2.8.2 *A universe of propositions*

We began this section by discussing the axiom of choice. While we were somewhat informal, if we wished to be more precise and have a single type denoting this axiom, we would need the ability to quantify over propositions within type theory. That is, we would require a type *of* propositions just as we considered a type of types in Section 2.6.

When considering a type of types, we ran into size issues (Girard’s paradox) which forced us to consider a hierarchy of universes U_0, U_1, \dots . Surprisingly, this hierarchy is *not* mandatory in the case of propositions. We may certainly take the predicative approach and add a hierarchy of universes of propositions paralleling U_i , but it is also consistent to a single *impredicative* universe which contains all propositions.

Remark 2.8.3. To give some intuition as to why this is not immediately contradictory, observe that no matter how we arrange the theory, there are multiple distinct propositions. Accordingly, the type of all propositions must contain distinct elements and so is not itself a proposition. We will never have the equivalent of $U : U$ for this reason alone. \diamond

Remark 2.8.4. The terms *predicative* and *impredicative* originate in philosophy and logic. Roughly, a construction is said to be impredicative if it may quantify over itself during the process of its construction while it is predicative if this is disallowed. In our situation, it will be possible to form codes within impredicative universe of propositions Ω which denote *e.g.*, $\Pi(\Omega, \dots)$ and therefore clearly “quantify over Ω ”. However, it is perhaps best to simply accept (im)predicative as purely technical terms, as their meaning in philosophy and logic is neither entirely precise nor consistent. \diamond

Both approaches have pros and cons. A single impredicative universe is a powerful extension to type theory and it is essentially mandatory if one wishes to formalize certain areas of mathematics (for instance, point-set topology) in type theory. On the other hand,

it substantially complicates the theory and a good deal of metatheory of type theory (Chapter 3) becomes exponentially more difficult. Accordingly, we shall detail both here.

Remark 2.8.5. Both predicative and impredicative universes of propositions appear in real-world proof assistants. For instance, Agda [Agda] features a predicative hierarchy of universes of propositions while Coq [Coq] has an impredicative universe. \diamond

We begin with the simpler case of a predicative universe of propositions. Our goal is to define a collection of universes (see Notation 2.6.4) $\Gamma \vdash \Omega_i$ type and $\Gamma.\Omega_i \vdash \mathbf{Prf}$ type which contains the subset of codes in U_i which decode to propositions. In fact, we have no need to add additional rules to our type theory to obtain such universes. We can already define them from U_i :

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \Omega_i := \Sigma(U_i, \text{isProp}_{\text{El}(A)}) \text{ type}} \qquad \frac{\vdash \Gamma \text{ cx} \quad \Gamma \vdash c : \Omega_i}{\Gamma \vdash \mathbf{Prf}_i(c) := \text{El}_i(\text{fst}(c)) \text{ type}}$$

Notation 2.8.6. We have chosen the notation $\mathbf{Prf}(\phi)$ as elements of this type should be viewed as **proofs** of ϕ .

This accomplishes our goal of collecting all propositions (of a certain size) into a single type without any further effort. Revisiting our goal of expressing the axiom of choice within type theory, we are now able to quantify over relations by quantifying over the type $A \times B \rightarrow \Omega_i$. However, we immediately become stuck once more: we do not have any proposition corresponding to \exists . After all, we have argued that Σ does not realize this notion and, while the argument is far from straightforward, one can argue that there is no suitable type in type theory.¹⁵

Of course, the absence of \exists is hardly the fault of Ω_i . If we wished, we could simply extend our theory with a type which behaves like \exists and march on. Remarkably, however, if we were to extend our theory with an impredicative universe of propositions \exists will become definable!

2.8.3 An impredicative universe of propositions

Let us now turn to an *impredicative* universe of propositions. In broad strokes, this universe will collapse together all the universes Ω_i into a single universe (Ω, \mathbf{Prf}) which contains all propositions. This requires a genuine extension to type theory and can no longer be achieved by combining existing connectives. As ever when adding new connectives to type theory, we should be mindful of Slogan 2.5.3. For expediency, we will begin by specifying Ω through inference rules and defer the statement of its mapping-in property to Exercise 2.44.

¹⁵The authors are grateful to Andrew Swan for explaining a proof of this fact.

We begin with the formation and elimination rules:

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \Omega \text{ type}} \quad \frac{\Gamma \vdash \phi : \Omega}{\Gamma \vdash \mathbf{Prf}(\phi) \text{ type}} \quad \frac{\Gamma \vdash \phi : \Omega \quad \Gamma \vdash x_0, x_1 : \mathbf{Prf}(\phi)}{\Gamma \vdash x_0 = x_1 : \mathbf{Prf}(\phi)}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash \phi : \Omega}{\Delta \vdash \mathbf{Prf}(\phi)[\gamma] = \mathbf{Prf}(\phi[\gamma]) \text{ type}} \quad \frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \Omega[\gamma] = \Omega \text{ type}}$$

In other words, (Ω, \mathbf{Prf}) forms some universe of propositions. It remains, however, to populate this universe. We shall do so with a single and exceptionally strong rule, reminiscent of the inconsistent `code` operation we discussed in Section 2.6:

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma.A.A[\mathbf{p}] \vdash \mathbf{q}[\mathbf{p}] = \mathbf{q} : A[\mathbf{p}^2]}{\Gamma \vdash \mathbf{code}(A) : \Omega}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type} \quad \Gamma.A.A[\mathbf{p}] \vdash \mathbf{q}[\mathbf{p}] = \mathbf{q} : A[\mathbf{p}^2]}{\Delta \vdash \mathbf{code}(A)[\gamma] = \mathbf{code}(A[\gamma]) : \Omega}$$

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma.A.A[\mathbf{p}] \vdash \mathbf{q}[\mathbf{p}] = \mathbf{q} : A[\mathbf{p}^2]}{\Gamma \vdash \mathbf{Prf}(\mathbf{code}(A)) = A \text{ type}} \quad \frac{\Gamma \vdash \phi : \Omega}{\Gamma \vdash \phi = \mathbf{code}(\mathbf{Prf}(\phi)) : \Omega}$$

We can equivalently specify this type using a mapping-in property:

Exercise 2.44. Construct a natural isomorphism between $\mathbf{Tm}(\Gamma, \Omega)$ and the following set:

$$\{\Gamma \vdash A \text{ type} \mid A \text{ is a proposition}\}$$

Conclude that Ω can be specified using Slogan 2.4.4.

Advanced Remark 2.8.7. Ω is *almost* a subobject classifier. The only obstruction is the equality of elements on Ω : if we wish for it to be a subobject classifier we could add an additional principle stating that logically equivalent propositions have equal codes in Ω . We shall return to this point in Section 5.2 under the name *propositional univalence*. \diamond

Remark 2.8.8. We have chosen a very strong formulation of Ω as it is the simplest to specify. Often one will encounter an impredicative universe of propositions which is essentially just closed under \forall . This is what is done in *e.g.*, `Coq [Coq]`. \diamond

We are now almost ready to deliver on our promise that the addition of Ω allows us to construct $\Gamma \vdash \mathbf{exists}(A, \phi) : \Omega$ for any $\Gamma \vdash A \text{ type}$ and $\Gamma.A \vdash \phi : \Omega$. Prior to doing so, however, we must explain what rules we expect it to satisfy. We therefore take a slight detour in order to discuss an adaptation of Slogan 2.5.3 to apply to propositions as well.

Interlude: mapping in and mapping out for propositions Previously, we have answered questions about what rules ought to govern a type using Slogan 2.5.3. This same principle can be used to give a systematic account of what rules ought to govern $\mathbf{Prf}(\mathbf{exists}(c, \phi))$ along with other propositions.

The crucial change is that while types are determined by mapping into or mapping out of other types, propositions are determined by their relation to other propositions. For instance, we may specify **exists** by means of the following mapping out property:

1. $\Gamma \vdash \mathbf{Prf}_i(\mathbf{exists}(c, \phi))$ type is a proposition
2. If $\Gamma \vdash A$ type is a proposition then there should be natural isomorphism of the following shape:

$$\mathbf{Tm}(\Gamma.\mathbf{Prf}_i(\mathbf{exists}(c, \phi)), A[\mathbf{p}]) \cong \mathbf{Tm}(\Gamma.\mathbf{El}_i(c).\mathbf{Prf}_i(\phi), A[\mathbf{p}^2])$$

Several points are worth emphasizing about this specification. First, unlike other mapping out properties we have seen in Section 2.5, $\Gamma \vdash A$ type is not allowed to depend on $\mathbf{Prf}_i(\mathbf{exists}(c, \phi))$. This is not the weakness it might first appear: $\mathbf{Prf}_i(\mathbf{exists}(c, \phi))$ is a proposition and so all its inhabitants are equal. It is even possible to derive the more complex version of the mapping-out property where $\Gamma.\mathbf{Prf}_i(\mathbf{exists}(c, \phi)) \vdash A$ type from the above principle using e.g., Σ . Accordingly, we have opted for the simpler version given above for clarity.

Second, it is perhaps surprising that **exists** intuitively parallels Σ and yet has a mapping-out rather than a mapping-in property. This is a consequence of the fact that Σ can be specified by *either* a mapping-in or mapping-out property. We opted for the former as it is more convenient, but we could have just as well specified it by a mapping-out property. On the other hand, **exists** does not enjoy this dual-character and has only the mapping-out property that we have specified above. See the final paragraph of Remark 2.5.1 for a more technical discussion of this point.

Should we add an exercise or something about this symmetry for sigma? Might go well in 2.5.4?

We can turn these properties into a pair of rules governing **exists** just as was done in Sections 2.4 and 2.5. By doing so, we obtain the following pair of rules:

$$\frac{\Gamma \vdash c : \mathbf{U}_i \quad \Gamma.\mathbf{El}_i(c) \vdash \phi : \mathbf{\Omega}_i \quad \Gamma \vdash a : \mathbf{El}_i(c) \quad \Gamma \vdash x : \mathbf{Prf}_i(\phi[\mathbf{id}.a] \quad \Gamma \vdash \mathbf{ex}(a, x) : \mathbf{Prf}_i(\mathbf{exists}(c, \phi))$$

$$\frac{\Gamma \vdash c : \mathbf{U}_i \quad \Gamma.\mathbf{El}_i(c) \vdash \phi : \mathbf{\Omega}_i \quad \Gamma \vdash A \text{ type} \quad \Gamma.A.A[\mathbf{p}] \vdash \mathbf{q}[\mathbf{p}] = \mathbf{q} : A[\mathbf{p}^2] \quad \Gamma.\mathbf{El}_i(c).\mathbf{Prf}_i(\phi) \vdash a : A[\mathbf{p}^2] \quad \Gamma \vdash x : \mathbf{Prf}_i(\mathbf{exists}(c, \phi))}{\Gamma \vdash \mathbf{split}(x, a) : A}$$

We note that there is no need to formulate either naturality, β , or η laws for these rules: all of them deal with elements of propositions and thus all possible equations governing them are automatically satisfied.

Remark 2.8.9. These rules are more-or-less translations of the familiar principles from higher-order logic:

$$\frac{\Gamma \mid \exists x : \tau. \phi \vdash \psi}{\Gamma, x : \tau \mid \phi \vdash \psi} \quad \frac{\Gamma, x : \tau \mid \phi \vdash \psi}{\Gamma \mid \exists x : \tau. \phi \vdash \psi}$$

This gives us some confidence that our specification faithfully encodes existential quantification. \diamond

This story can be played out to specify all propositional versions of all the connectives of type theory. This done, for instance, by Gilbert et al. [Gil+19].

Defining exists using Ω

Notation 2.8.10. For the remainder of this section, we shall use informal notation for type theory to increase readability.

We now (finally) deliver on our promise.

Lemma 2.8.11. *Given a type A and $\phi : A \rightarrow \Omega$, there exists an element $\exists a : A. \phi a : \Omega$ along with an isomorphism between the following types for all propositions B :*

$$\mathbf{Prf}(\exists a : A. \phi a) \rightarrow B \cong ((a : A) \rightarrow \mathbf{Prf}(\phi a)) \rightarrow B[\mathbf{p}^2]$$

Remark 2.8.12. This proof builds on the same ideas as classical Church-encodings from the untyped lambda calculus or System F: using quantification over Ω , we may define an element of Ω directly from its mapping-out property. \diamond

Proof. We shall define \exists such that it satisfies the desired isomorphism more-or-less by definition. Explicitly, consider the following type:

$$X = (\psi : \Omega)((a : A)(\mathbf{Prf}(\phi a)) \rightarrow \mathbf{Prf}(\psi)) \rightarrow \mathbf{Prf}(\psi)$$

By virtue of Exercise 2.42, X is a proposition: it is a dependent product whose target is a proposition. Accordingly, we may form $\exists a : A. \phi a = \mathbf{code}(X)$. It remains to show that this satisfies the expected isomorphism for all propositions B :

$$(\mathbf{Prf}(\mathbf{exists}(A, \phi)) \rightarrow B) \cong ((a : A)(\mathbf{Prf}(\phi a)) \rightarrow B)$$

We can replace B in the above with $\mathbf{Prf}(\xi)$ where $\xi = \mathbf{code}(B)$. In this case, we are left to show that the following two types are isomorphic:

$$Z_0 = ((\psi : \Omega)((a : A)(\mathbf{Prf}(\phi a)) \rightarrow \mathbf{Prf}(\psi)) \rightarrow \mathbf{Prf}(\psi)) \rightarrow \mathbf{Prf}(\xi)$$

$$Z_1 = (a : A)(\mathbf{Prf}(\phi a)) \rightarrow \mathbf{Prf}(\psi)$$

We construct the necessary functions as follows:

$$f : Z_0 \rightarrow Z_1$$

$$f z_0 a x = z_0(\lambda \psi k. k a x)$$

$$g : Z_1 \rightarrow Z_0$$

$$f z_1 k = k \xi z_1$$

These functions are automatically inverse to each other: both Z_0 and Z_1 are propositions and so any function $Z_i \rightarrow Z_i$ is the identity function. \square

Lemma 2.8.13. *Given two propositions $\phi, \psi : \Omega$, there exists a proposition $\phi \vee \psi : \Omega$ such that the following isomorphism holds for all propositions B :*

$$(\mathbf{Prf}(\phi \vee \psi) \rightarrow B) \cong (\mathbf{Prf}(\phi) \rightarrow B) \times (\mathbf{Prf}(\psi) \rightarrow B)$$

Proof. We define $\phi \vee \psi$ to be $\exists b : \mathbf{Bool}. \mathbf{if}(b, \phi, \psi)$. All desired properties then follow from Lemma 2.8.11. \square

Exercise 2.45. Give a direct construction of $\phi \vee \psi$ using the same techniques to construct $\exists a : A. \phi a$.

Remark 2.8.14. The reader might guess that $\mathbf{Prf}(\phi \vee \psi)$ and $\mathbf{Prf}(\phi) + \mathbf{Prf}(\psi)$ never coincide. In fact, however, they are equivalent in one important situation: when ϕ and ψ are disjoint propositions. That is, when $\mathbf{Prf}(\phi) \times \mathbf{Prf}(\psi) \rightarrow \mathbf{Void}$. \diamond

We shall end our discussion of universes of propositions here. In Section 5.1, we shall resume this thread in the context of *propositional univalence*.

2.8.4 Logic in type theory

Finally, we have developed the tools necessary to state the “proper” axiom of choice within type theory and we finally ask the obvious question: does this principle hold in type theory? We can phrase this question in a more general manner: what is sort of logic do propositions in type theory support?

We shall not keep the reader in suspense: the answer is *constructive* logic. That is, logic without the axiom $\phi \vee \neg \phi$, the axiom of choice, or similar principles. We emphasize that we do not refute these principles: it is perfectly consistent to postulate choice within type theory. So while type theory does not validate the axiom of choice, neither does it prove

any of the specialized “anti-classical” reasoning principles often contemplated by logicians. More technically, one can prove the following result (though its proof is out-of-scope for this book).

Theorem 2.8.15. *Neither of the following types are inhabited within type theory:*

1. $(\phi : \Omega) \rightarrow \mathbf{Prf}(\mathbf{Prf}(\phi) \vee \mathbf{code}(\phi \rightarrow \mathbf{Void}))$
2. $(\phi : \Omega) \rightarrow \mathbf{Prf}(\mathbf{Prf}(\phi) \vee \mathbf{code}(\phi \rightarrow \mathbf{Void})) \rightarrow \mathbf{Void}$

In other words, we can neither prove nor refute ϕ or $\neg\phi$ for all ϕ .

A similar result is available for the axiom of choice.

Theorem 2.8.16. *Let us denote by LogicalChoice the following element of Ω :*

$$\mathbf{code}(((a : A) \rightarrow \mathbf{Prf}(\exists b : B. P(a, b))) \rightarrow \mathbf{Prf}(\exists f : A \rightarrow B. (a : A) \rightarrow \mathbf{Prf}(P(a, f(a))))))$$

Neither of the following types are inhabited within type theory:

1. $\mathbf{Prf}(\text{LogicalChoice})$
2. $\mathbf{Prf}(\text{LogicalChoice}) \rightarrow \mathbf{Void}$

Perhaps more surprisingly, type theory is not only agnostic to the validity of classical statements like the law of excluded middle and choice. It is also consistent with principles which are classically *false*. For example, type theory is consistent with “Church’s law”:

$$(f : \mathbf{Nat} \rightarrow \mathbf{Nat}) \rightarrow \exists n : \mathbf{Nat}. \text{“}n \text{ is the code of a Turing machine computing } f\text{”}$$

To see that this contradicts LEM, observe that with LEM we may write a function stating sending a Turing machine to a boolean indicating whether or not it halts. Such a function can never be tracked by another Turing machine.

Various other anticlassical principles are also consistent *e.g.*, the statement that all functions from the real numbers to the real numbers are consistent. This flexibility can be used to shape type theory into a powerful domain-specific language for a particular sort of reasoning. This approach to mathematics is often referred to *synthetic mathematics*. While we shall not explore it further in this book, in Chapter 5 we will explore a particular application of this philosophy, namely homotopy type theory.

Further reading

The literature on type theory is unfortunately neither notationally nor conceptually coherent, particularly regarding syntax and how it is defined. We summarize a number of important references that most closely match the perspective outlined in this book; note however that many references will agree in some ways and differ in others.

Historical references Nearly all of the ideas in this chapter can be traced back in some form to the philosopher Per Martin-Löf, whose collected works are available in the GitHub repository [michaelt/martin-lof](https://github.com/michaelt/martin-lof). Over the decades, Martin-Löf has considered many different variations on type theory; the closest to our presentation are his notes on substitution calculus [ML92] and the “Bibliopolis book” presenting what is now called extensional type theory [ML84b]. For a detailed philosophical exploration of the *judgmental methodology* that types internalize judgmental structure, see his “Siena lectures” [ML96]. Finally, the book *Programming in Martin-Löf’s Type Theory* [NPS90] remains one of the best pedagogical introductions to type theory as formulated in Martin-Löf’s logical framework.

Syntax of dependent type theory The presentation of type theory most closely aligned to ours can be found in the second author’s Ph.D. thesis [Gra23, Chapter 2]. Another valuable reference is Hofmann’s *Syntax and Semantics of Dependent Types* [Hof97, Sections 1 and 2], which as the title suggests, presents the syntax of type theory and connects it to semantical interpretations. Hofmann is very careful in his definition of syntax, but the technical details of capture-avoiding substitution and presyntax have largely been supplanted by subsequent work on logical frameworks, so we suggest that readers gloss over these technical details.

Categorical semantics The book *Categories for Types* [Cro94] is a gentle introduction to the categorical semantics of the simply-typed lambda calculus and related theories; Castellan, Clairambault, and Dybjer [CCD21] discuss how to scale from such models to categories with families [Dyb96], the categorical counterpart of the substitution calculus. Readers can consult Hofmann [Hof97] for concrete examples of categories with families. Finally, we recommend Awodey’s paper on *natural models* [Awo18] for a more categorically-natural formulation of categories with families, as well as an excellent description of the *local universes* strictification procedure for producing models of dependent type theory from categories with enough structure [LW15].

Logical frameworks In this book we have attempted to largely sidestep the question of what constitutes a valid collection of inference rules. The mathematics of syntax can and has occupied entire books, but in short, the natural families of constants and

isomorphisms considered in this chapter can be formulated precisely in systems known as *logical frameworks*. A good introduction to logical frameworks is the seminal work of Harper, Honsell, and Plotkin [HHP93] on the Edinburgh Logical Framework, in which object-level judgments can be encoded as meta-level types.

For logical frameworks better suited to defining dependent type theory in particular, we refer readers to the *generalized algebraic theories* of [Car86] (or the tutorial on this subject by Sterling [Ste19]), or to *quotient inductive-inductive types* [AK16; Dij17; KKA19; Kov22]. For logical frameworks specifically designed to accommodate the binding and substitution of dependent type theory, we refer the reader to the Ph.D. theses of Haselwarter [Has21] and Uemura [Uem21].

Metatheory and implementation

In Chapter 2 we carefully defined Martin-Löf type theory as a formal mathematical object: a kind of “algebra” of indexed sets (of types and terms) equipped with various operations. We believe this perspective is essential to understanding both the *what* and the *why* of type theory, providing both a precise definition that can be unfolded into inference rules, as well as an explanation of what these rules intend to axiomatize.

This perspective is not, however, how most users of type theory interact with it. Most users of type theory interact with *proof assistants*, software systems for interactively developing and verifying large-scale proofs in type theory. Even when type theorists work on paper rather than on a computer, many of the conveniences of proof assistants bleed into their informal notation. Indeed, in Chapter 1 we used definitions, implicit arguments, data type declarations, and pattern matching without a second thought.

Although this book focuses on theoretical rather than practical considerations, it is impossible to discuss the design space of type theory without discussing the pragmatics of proof assistants, as these have exerted a profound influence on the theory. Our goal in this chapter is to explain how to square our mathematical notion of type theory with (idealized) implementations¹ of type theory, and to discover and unpack the substantial constraints that the latter must place on the former.

In this chapter In Section 3.1 we axiomatize the core functionality of proof assistants in terms of algorithmic elaboration judgments, and outline a basic implementation. In Section 3.2 we continue to refine our implementation, taking a closer look at how the equality judgments of type theory impact elaboration, and the metatheoretic properties we need equality to satisfy. In Section 3.3 we consider how to extend our elaborator to account for definitions. In Section 3.4 we discuss other metatheorems of type theory and their relationship to program extraction. In Section 3.5 we construct a set-theoretic model of extensional type theory and explore some of its metatheoretic consequences—including a counterexample to one of the properties discussed in Section 3.2. Finally, in Section 3.6 we disprove a second important metatheoretic property, leading us to consider alternatives to extensional type theory (Chapter 2) in Chapters 4 and 5.

Goals of the chapter By the end of this chapter, you will be able to:

- Explain why and how we define type-checking in terms of elaboration.

¹At the end of this chapter, we provide some pointers to literature and implementations specifically geared to readers interested in learning how to actually implement type theory.

- Define the consistency, canonicity, normalization, and invertibility metatheorems, and identify why each is important.
- Explain which metatheorems are disrupted by extensional equality, and sketch why.

3.1 *A judgmental reconstruction of proof assistants*

What exactly is the relation between Agda code (or the code in Chapter 1) and the type theory in Chapter 2? Certainly, Coq and Agda—even without extensions—include many convenience features that the reader would not be surprised to see omitted in a theoretical description of type theory: implicit arguments, typeclasses/instance arguments, libraries, reflection, tactics... For the moment we set aside not only these but even more fundamental features such as data type declarations, pattern matching, and the ability to write definitions, in order to consider the simplest possible “Agda”: a *type-checker*. That is, our idealized Agda takes as input two expressions e and τ and *accepts* in the case that e is a closed term of closed type τ , and *rejects* if not.

Slogan 3.1.1. *Proof assistants are fancy type-checkers.*

Remark 3.1.2. For the purposes of this book, “proof assistant” refers only to proof assistants in the style of Coq, Agda, and Lean. In particular, we will not discuss LCF-style systems [GMW79] such as Nuprl [Con+85] and Andromeda [Bau+21], or systems not based on dependent type theory, such as Isabelle [NPW02] or HOL Light [Har09]. \diamond

Convenience features of proof assistants are generally aimed at making it easier for users to write down the inputs e and τ , perhaps by allowing some information to be omitted and reconstructed mechanically, or even by presenting a totally different interface for building e and τ interactively or from high-level descriptions. We start our investigation with the most generous possible assumptions—in which e and τ contain all the information we might possibly need, including type annotations—and will find that type-checking is already a startlingly complex problem.

Remark 3.1.3. The title of this section is an homage to *A judgmental reconstruction of modal logic* [PD01], an influential article that reconsiders intuitionistic modal logic under the mindset that *types internalize judgmental structure*. \diamond

3.1.1 *Type-checking as elaboration*

In Section 2.1 we emphasized that we do *not* assume that the types and terms of type theory are obtained as the “well-formed” subsets of some collections of possibly-ill-formed

$$\begin{array}{ll}
\text{Pretypes} & \tau := (\text{Pi } \tau \tau) \mid (\text{Sigma } \tau \tau) \mid \text{Unit} \mid \text{Uni} \mid (\text{El } e) \mid \dots \\
\text{Preterms} & e := (\text{var } i) \mid (\text{lam } \tau \tau e) \mid (\text{app } \tau \tau e e) \mid (\text{pair } \tau \tau e e) \mid (\text{fst } \tau \tau e) \mid \dots \\
\text{Indices} & i := 0 \mid 1 \mid 2 \mid \dots
\end{array}$$

Figure 3.1: Syntax of pretypes and preterms.

pretypes or *preterms*, nor do we even assume that they are obtained as “ $\beta\eta$ -equivalence classes” of well-formed-but-unquotiented terms.

Instead, *types* and *terms* are just the elements of the sets $\text{Ty}(\Gamma)$ and $\text{Tm}(\Gamma, A)$, which are defined in terms of each other and the sets Cx and $\text{Sb}(\Delta, \Gamma)$. When we write e.g. $\lambda(b)$, we are naming a particular element of a particular set $\text{Tm}(\Gamma, \Pi(A, B))$ obtained by applying the “ Π -introduction” map to $b \in \text{Tm}(\Gamma.A, B)$; in particular, the values of Γ, A, B should be regarded as implicitly present, as they are in Appendix A where we write $\lambda_{\Gamma, A, B}(b)$.

In Chapter 2 we reaped the benefits of this perspective, but it has come time to pay the piper: what, then, is a type-checker supposed to take as input? We certainly cannot say that a type-checker is given “a type A and a term a ” because this assumes that A and a are well-formed. Type-checking *cannot be a membership query*; instead, it is a *partial function* from concrete syntax to the sets of genuine types and terms. For an input expression to “type-check” means that it *names* a type/term, not that it “is” one (which is a meta-type error, as types/terms are mathematical objects, and input expressions are strings).

For simplicity we assume that the inputs to type-checkers are not strings but abstract syntax trees (or well-formed formulas) conforming to the simple grammar in Figure 3.1.² We call these semi-structured input expressions *pretypes* τ and *preterms* e , and write them as teletype s-expressions. In programming language theory, the process of mapping semi-structured input expressions into structured core language terms is known as *elaboration*.

Slogan 3.1.4. *Type-checkers for dependent type theory are elaborators.*

Remark 3.1.5. What is the relationship between features of the concrete syntax of a proof assistant, and features of the core syntax? According to Slogan 3.1.4, the concrete syntax should be seen as “instructions” for building core syntax. These instructions may be very close to or very far from that core syntax, but in either case, new user-facing features should only induce new core primitives when they cannot be (relatively compositionally) accounted for by the existing core language. \diamond

Algorithmic judgments Elaborators are partial functions that recursively consume pretypes and preterms (abstract syntax trees) and produce types and terms. In a real proof assistant, types and terms are of course not abstract mathematical entities but elements of

²In other words, we only consider input expressions that successfully parse; expressions that fail to parse (e.g., because their parentheses are mismatched) automatically fail to type-check.

some data type, but for our purposes we will imagine an idealized elaborator that outputs elements of $\text{Ty}(\Gamma)$ and $\text{Tm}(\Gamma, A)$. We present this elaborator not as functional programs written in pseudocode, but as *algorithmic judgments* defined by inference rules. Unlike the rules in Chapter 2, these rules are intended to define an algorithm, so we will take care to ensure that any given elaboration judgment can be derived by at most one rule. (In other words, we define our elaborator as a deterministic logic program.)

We have already argued that pretype elaboration should take as input a pretype τ and output a type A , but what about contexts? Just as well-formedness of closed types ($1 \vdash \Pi(A, B)$ type) refers to well-formedness of open types ($1.A \vdash B$ type), it is perhaps unsurprising that elaborating closed pretypes requires elaborating open pretypes. However, we note that we do not need or want “precontexts”; we will only descend under binders after successfully elaborating their pretypes. For example, to elaborate $(\text{Pi } \tau_0 \tau_1)$ we will first elaborate τ_0 to the closed type A , and only then in context $1.A$ elaborate τ_1 to B .

Thus our two main algorithmic elaboration judgments are as follows:

1. $\Gamma \vdash \tau \text{ type} \rightsquigarrow A$ asserts that elaborating the pretype τ relative to $\vdash \Gamma \text{ cx}$ succeeds and produces the type $\Gamma \vdash A \text{ type}$.
2. $\Gamma \vdash e : A \rightsquigarrow a$ asserts that elaborating the preterm e relative to $\vdash \Gamma \text{ cx}$ and $\Gamma \vdash A \text{ type}$ succeeds and produces the term $\Gamma \vdash a : A$.

In pseudocode, the first judgment corresponds to a partial function $\text{elabTy}(\Gamma, \tau) = A$ with the invariant that if $\vdash \Gamma \text{ cx}$ and elabTy terminates successfully, then $\Gamma \vdash A \text{ type}$. Likewise, the second judgment is a partial function $\text{elabTm}(\Gamma, A, e) = a$ whose successful outputs are terms $\Gamma \vdash a : A$.

Elaborating pretypes The rules for $\Gamma \vdash \tau \text{ type} \rightsquigarrow A$ are straightforward translations of the type-well-formedness rules of Chapter 2. (When it is necessary to contrast algorithmic and non-algorithmic rules, the latter are often referred to as *declarative*.)

$$\frac{\Gamma \vdash \tau_0 \text{ type} \rightsquigarrow A \quad \Gamma.A \vdash \tau_1 \text{ type} \rightsquigarrow B}{\Gamma \vdash (\text{Pi } \tau_0 \tau_1) \text{ type} \rightsquigarrow \Pi(A, B)} \quad \frac{\Gamma \vdash \tau_0 \text{ type} \rightsquigarrow A \quad \Gamma.A \vdash \tau_1 \text{ type} \rightsquigarrow B}{\Gamma \vdash (\text{Sigma } \tau_0 \tau_1) \text{ type} \rightsquigarrow \Sigma(A, B)}$$

$$\frac{}{\Gamma \vdash \text{Unit type} \rightsquigarrow \mathbf{Unit}} \quad \frac{}{\Gamma \vdash \text{Uni type} \rightsquigarrow \mathbf{U}} \quad \frac{\Gamma \vdash e : \mathbf{U} \rightsquigarrow a}{\Gamma \vdash (\text{El } e) \text{ type} \rightsquigarrow \mathbf{El}(a)}$$

3.1.2 Elaborating preterms: the problem of type equality

Elaborating preterms is significantly more fraught. But first, let us remind ourselves of the process of type-checking $(\text{lam } \tau_0 \tau_1 e) : \tau$. First, we attempt to elaborate the

pretype $1 \vdash \tau$ type $\rightsquigarrow C$; if this succeeds, we then attempt to elaborate the preterm $1 \vdash (\text{lam } \tau_0 \ \tau_1 \ e) : C \rightsquigarrow c$. If this also succeeds, then the type-checker reports success, having transformed the input presyntax to a well-formed term $1 \vdash c : C$.

Since `lam` is our presyntax for λ , elaborating `lam` via $1 \vdash (\text{lam } \tau_0 \ \tau_1 \ e) : C \rightsquigarrow c$ should produce a term $c := \lambda_{1,A,B}(b)$ for some A, B, b determined by τ_0, τ_1, e respectively. We determine these by a series of recursive calls to the elaborator: first $\Gamma \vdash \tau_0$ type $\rightsquigarrow A$, then $\Gamma.A \vdash \tau_1$ type $\rightsquigarrow B$, and finally $\Gamma.A \vdash e : B \rightsquigarrow b$. Note that these steps must be performed sequentially and in this order, because each step uses the outputs of the previous steps as inputs: we elaborate τ_1 in a context extended by A , the result of elaborating τ_0 , and we elaborate e at type B , the result of elaborating τ_1 .

At the end we obtain $\Gamma.A \vdash b : B$, and thence by Π -introduction a term $1 \vdash \lambda_{1,A,B}(b) : \Pi_1(A, B)$ that should be the elaborated form of e . But the elaborated form of e is supposed to have type C —the result of elaborating τ ! Thus before returning $\lambda_{1,A,B}(b)$ we need to check that $1 \vdash C = \Pi(A, B)$ type. This is where “type-checking” actually happens: we have seen that τ determines a real type and that e determines a real term, but until this point we have not actually checked whether “ e has type τ ”.

In pseudocode, we can define elaboration of $(\text{lam } \tau_0 \ \tau_1 \ e)$ as follows:

```

elabTm( $\Gamma, C, (\text{lam } \tau_0 \ \tau_1 \ e)$ ) =
  let  $A = \text{elabTy}(\Gamma, \tau_0)$  in
  let  $B = \text{elabTy}(\Gamma.A, \tau_1)$  in
  let  $b = \text{elabTm}(\Gamma.A, B, e)$  in
  if  $(\Gamma \vdash C = \Pi_\Gamma(A, B)$  type) then return  $\lambda_{\Gamma,A,B}(b)$  else error

```

or equivalently, in algorithmic judgment notation:

$$\frac{\Gamma \vdash \tau_0 \text{ type } \rightsquigarrow A \quad \Gamma.A \vdash \tau_1 \text{ type } \rightsquigarrow B \quad \Gamma.A \vdash e : B \rightsquigarrow b \quad \Gamma \vdash C = \Pi(A, B) \text{ type}}{\Gamma \vdash (\text{lam } \tau_0 \ \tau_1 \ e) : C \rightsquigarrow \lambda_{\Gamma,A,B}(b)}$$

This will be the only rule that concludes $\Gamma \vdash e : C \rightsquigarrow c$ for $e := (\text{lam } \tau_0 \ \tau_1 \ e)$, ensuring that this rule “is the `lam` clause of `elabTm`,” so to speak. Elaboration of other introduction forms will follow a similar pattern.

Exercise 3.1. Write the algorithmic rule for elaborating the preterm $(\text{pair } \tau_0 \ \tau_1 \ e_0 \ e_1)$.

Let us pause to make several remarks. First, note that our algorithm needs to check judgmental equality of types $\Gamma \vdash C = \Pi_\Gamma(A, B)$ type. This step is, at least implicitly, part of all type-checking algorithms for all programming languages: if we define a function of type $A \rightarrow B$ that returns e , we have to check whether the type of e matches the declared return type B . Sometimes this is as simple as checking the syntactic equality of two type expressions, but often this is non-trivial, perhaps a subtyping check.

In our present setting, checking type equality is *extremely* non-trivial. Suppose that $C := \mathbf{El}(c)$ and so we are checking $\Gamma \vdash \mathbf{El}(c) = \mathbf{\Pi}(A, B)$ type for $\Gamma \vdash c : U$. This type equality depends on the entire equational theory of *terms*: we may need to “rewrite along” arbitrarily many term equations before concluding $\Gamma \vdash c = \mathbf{pi}(c_0, c_1) : U$; this only reduces the problem to $\Gamma \vdash \mathbf{\Pi}(\mathbf{El}(c_0), \mathbf{El}(c_1)) = \mathbf{\Pi}(A, B)$ type for which it suffices to check $\Gamma \vdash \mathbf{El}(c_0) = A$ type and $\Gamma.A \vdash \mathbf{El}(c_1) = B$ type, each of which may once again require arbitrary amounts of computation. We will revisit this point in Section 3.2.1.

Secondly, note that we have assumed for now that the preterm $(\mathbf{lam} \tau_0 \tau_1 e)$ contains pretype annotations τ_0, τ_1 telling us the domain and codomain of the $\mathbf{\Pi}$ -type. In practice, a type-checker is essentially unusable unless it can *reconstruct* (most of) these annotations; we describe this reconstruction process in Section 3.2.2.

Remark 3.1.6. Naïvely, one might think that *including* these annotations is the source of our problem, because it forces us to compare the type C computed from τ to the type $\mathbf{\Pi}(A, B)$ computed from the annotations τ_0, τ_1 . This is not the case. If we omit τ_0, τ_1 , then to elaborate e we must *recover* A and B from C , which upgrades “does $\Gamma \vdash C = \mathbf{\Pi}(A, B)$ type?” to the strictly harder question “do there *exist* A, B such that $\Gamma \vdash C = \mathbf{\Pi}(A, B)$ type?” In addition, we will need to wonder whether this existence is unique: otherwise, it could be that $\Gamma.A \vdash e : B \rightsquigarrow b$ for some choices of A, B but not others. \diamond

Elaborating elimination forms is not much harder than elaborating introduction forms. To elaborate $(\mathbf{app} \tau_0 \tau_1 e_0 e_1)$, we elaborate the pretype annotations $\Gamma \vdash \tau_0$ type $\rightsquigarrow A$ and $\Gamma.A \vdash \tau_1$ type $\rightsquigarrow B$ in sequence, then the function $\Gamma \vdash e_0 : \mathbf{\Pi}(A, B) \rightsquigarrow f$ and its argument $\Gamma \vdash e_1 : A \rightsquigarrow a$ in either order, before finally checking that the type of the computed term $\mathbf{app}_{\Gamma, A, B}(f, a)$, namely $B[\mathbf{id}.a]$, agrees with the expected type C .

$$\frac{\Gamma \vdash \tau_0 \text{ type } \rightsquigarrow A \quad \Gamma.A \vdash \tau_1 \text{ type } \rightsquigarrow B \quad \Gamma \vdash e_0 : \mathbf{\Pi}(A, B) \rightsquigarrow f \quad \Gamma \vdash e_1 : A \rightsquigarrow a \quad \Gamma \vdash C = B[\mathbf{id}.a] \text{ type}}{\Gamma \vdash (\mathbf{app} \tau_0 \tau_1 e_0 e_1) : C \rightsquigarrow \mathbf{app}_{\Gamma, A, B}(f, a)}$$

Elaboration of other elimination forms follows a similar pattern. The only remaining case is term variables $(\mathbf{var} i)$, which we have chosen to represent as de Bruijn indices. To elaborate $(\mathbf{var} i)$ we check that the context has length at least $i + 1$; if so, then it remains only to check that the type of the variable $\mathbf{q}[\mathbf{p}^i]$ agrees with the expected type.

$$\frac{\Gamma = \Gamma'.A_i.A_{i-1} \cdots .A_0 \quad \Gamma \vdash C = A_i[\mathbf{p}^{i+1}] \text{ type}}{\Gamma \vdash (\mathbf{var} i) : C \rightsquigarrow \mathbf{q}[\mathbf{p}^i]}$$

In the above rule, our algorithm needs to check judgmental equality of *contexts*, and to project Γ and A from $\Gamma.A$. Unlike for type equality, we have no rules generating non-trivial context equalities, so structural induction on contexts is perfectly well-defined.

Remark 3.1.7. It is straightforward to extend our concrete syntax to support named variables: in our elaboration judgments, we replace Γ with an *environment* Θ that is a list of pairs of genuine types with the “surface name” of the corresponding term variable. Every environment determines a context by forgetting the names; in the variable elaboration rule, we simply look up the de Bruijn index corresponding to the given name. \diamond

Exercise 3.2. Write the algorithmic rules for elaborating $(fst\ \tau_0\ \tau_1\ e)$ and $(snd\ \tau_0\ \tau_1\ e)$.

3.2 Metatheory for type-checking

In Section 3.1 we saw that we can reduce type-checking to the problem of deciding the equality of types (at least, assuming that our input preterms have all type annotations). Deciding the equality of types in turn requires deciding the equality of terms, particularly in the presence of universes (Section 2.6.2). It is far from obvious that these relations are decidable—in fact, as we will see in Section 3.6, they are actually *undecidable* for the theory presented in Chapter 2—and proving their decidability relies on a difficult metatheorem known as *normalization*. In this section, we continue our exploration of elaboration with an emphasis on normalization and other metatheorems necessary for type-checking.

Remark 3.2.1. Recall from Section 2.1 that a *metatheorem* is just an ordinary theorem in the ambient metatheory, particularly one concerning the object type theory. \diamond

Before we can discuss computability-theoretic properties of the judgments of type theory, however, we must fix an encoding. We have taken pains to treat the rules of type theory as defining abstract sets $Ty(\Gamma)$ and $Tm(\Gamma, A)$ equipped with functions (type and term formers) satisfying various equations (β and η laws), which is the right perspective for understanding the mathematical structure of type theory. But to discuss the *computational* properties of type theory it is essential to exhibit an effective encoding of types and terms that is suitable for manipulation by a Turing machine or other model of computation: Turing machines cannot take mathematical entities as inputs, and whether equality of types is decidable can depend on how we choose to encode them!

This is analogous to the issue that arises in elementary computability theory when formalizing the halting problem: we must agree on how to encode Turing machines as inputs to other Turing machines, and we must ensure that this encoding is suitably effective. It is possible to pick an encoding of computable functions that trivializes the halting problem, at the expense of this encoding itself necessarily being uncomputable.

Returning to type theory, derivation trees of inference rules (e.g., as in Appendix A) turn out to be a perfectly suitable encoding. That is, when discussing computability-theoretic properties of types, terms, and equality judgments, we shall assume that each of these is encoded by equivalence classes of closed derivation trees; for example, we encode $Ty(\Gamma)$

by the set of derivation trees with root $\Gamma \vdash A$ type for some A . (Just as there are many Turing machines realizing any given function $\mathbb{N} \rightarrow \mathbb{N}$, there will be many derivation trees encoding any given type $A \in \text{Ty}(\Gamma)$.) When we say “equality of types is decidable,” what we shall mean is that “it is decidable whether two derivations encode the same type.” But having fixed a convention, we will avoid belaboring the point any further.

3.2.1 Normalization and the decidability of equality

To complete the pretype and preterm elaboration algorithms presented in Section 3.1, it remains only to show that type and term equality are decidable, which is equivalent to the following normalization condition.

Remark 3.2.2. Type and term equality are automatically *semidecidable* because derivation trees are recursively enumerable. That is, to check whether two types $A, B \in \text{Ty}(\Gamma)$ are equal, we can enumerate every derivation tree of type theory, terminating if we encounter a derivation of $\Gamma \vdash A = B$ type. Obviously, this is not a realistic implementation strategy. \diamond

Definition 3.2.3. A *normalization structure* for a type theory is a pair of computable, injective functions $\text{nfTy} : \text{Ty}(\Gamma) \rightarrow \mathbb{N}$ and $\text{nfTm} : \text{Tm}(\Gamma, A) \rightarrow \mathbb{N}$.

Definition 3.2.4. A type theory enjoys *normalization* if it admits a normalization structure.

The reader may find these definitions surprising: where did \mathbb{N} come from, and where is the rest of the definition? We have chosen \mathbb{N} because it is a countable set with decidable equality, but any other such set would suffice. In practice, one instead defines two sets of abstract syntax trees TyNf , TmNf with discrete equality, and constructs a pair of computable, injective functions $\text{nfTy} : \text{Ty}(\Gamma) \rightarrow \text{TyNf}$ and $\text{nfTm} : \text{Tm}(\Gamma, A) \rightarrow \text{TmNf}$. It is trivial to exhibit computable, injective Gödel encodings of TyNf and TmNf , which when composed with nfTy , nfTm exhibit a normalization structure in the sense of Definition 3.2.3.

As for Definition 3.2.3 being sufficient, the force of normalization is that it gives us a decision procedure for type/term equality as follows: given $A, B \in \text{Ty}(\Gamma)$, A and B are equal if and only if $\text{nfTy}(A) = \text{nfTy}(B)$ in \mathbb{N} . Asking for these maps to be computable ensures that this procedure is computable; injectivity ensures that it is *complete* in the sense that $\text{nfTy}(A) = \text{nfTy}(B)$ implies $A = B$. The *soundness* of this procedure—that $A = B$ implies $\text{nfTy}(A) = \text{nfTy}(B)$ —is implicit in the statement that nfTy is a function out of $\text{Ty}(\Gamma)$, the set of types considered modulo judgmental equality.

Warning 3.2.5. In Section 3.6 we shall see that extensional type theory *does not* admit a normalization structure, but we will proceed under the assumption that the theory we are elaborating satisfies normalization. In Chapter 4 we will see how to modify our type theory to substantiate this assumption.

Assuming normalization, we can define algorithmic type and term equality judgments

1. $\Gamma \vdash A \Leftrightarrow B$ type asserts that the types $\Gamma \vdash A$ type and $\Gamma \vdash B$ type are judgmentally equal according to some decision procedure.
2. $\Gamma \vdash a \Leftrightarrow b : A$ asserts that the terms $\Gamma \vdash a : A$ and $\Gamma \vdash b : A$ are judgmentally equal according to some decision procedure.

as follows:

$$\frac{\text{nfTy}(A) = \text{nfTy}(B)}{\Gamma \vdash A \Leftrightarrow B \text{ type}} \qquad \frac{\text{nfTm}(a) = \text{nfTm}(b)}{\Gamma \vdash a \Leftrightarrow b : A}$$

We notate algorithmic equality differently from the declarative equality judgments $\Gamma \vdash A = B$ type and $\Gamma \vdash a = b : A$ to stress that their definitions are completely different, even though (by our argument above) two types/terms are algorithmically equal if and only if they are declaratively equal. We thus complete the elaborator from Section 3.1 by replacing the “calls” to $\Gamma \vdash C = \Pi(A, B)$ type with calls to $\Gamma \vdash C \Leftrightarrow \Pi(A, B)$ type.

Remark 3.2.6. It may seem surprising that normalization is so difficult; why can’t algorithmic equality just *orient* each declarative equality rule (e.g., $\text{fst}(\text{pair}(a, b)) \rightsquigarrow a$) and check whether the resulting rewriting system is confluent and terminating? Unfortunately, while this strategy suffices for some dependent type theories such as the calculus of constructions [CH88], it is very difficult to account for judgmental η rules. (What direction should $p \Leftrightarrow \text{pair}(\text{fst}(p), \text{snd}(p))$ go? What about the η rule of **Unit**, $a \Leftrightarrow \text{tt}$?) These rules require a type-sensitive decision procedure known as *normalization by evaluation*, whose soundness and completeness for declarative equality is nontrivial [ACD07; Abe13]. \diamond

Exercise 3.3. We argued that the existence of a normalization structure implies that judgmental equality is decidable. In fact, this is a biimplication. Assume that definitional equality is decidable, and construct from this a normalization structure. (Hint: some classical reasoning is required, such as Markov’s principle or the law of excluded middle.)

Exercise 3.4. We have sketched how to use normalization to obtain a type-checking algorithm. This, too, is a biimplication. Using Exercise 3.3, show that the ability to decide type-checking implies that normalization holds.

3.2.2 Injectivity and bidirectional type-checking

We have seen how to define a rudimentary elaborator for type theory assuming that normalization holds, but the preterms that we can elaborate (Figure 3.1) are quite verbose, making our proof assistant more of a proof adversary. For instance, function application ($\text{app } \tau_0 \tau_1 e_0 e_1$) requires annotations for both the domain and codomain of the Π -type.

$$\begin{aligned} \text{Pretypes } \tau &:= (\text{Pi } \tau \tau) \mid (\text{Sigma } \tau \tau) \mid \text{Unit} \mid \text{Uni} \mid (\text{El } e) \mid \dots \\ \text{Preterms } e &:= (\text{var } i) \mid (\text{chk } e \tau) \mid (\text{lam } e) \mid (\text{app } e e) \mid (\text{pair } e e) \mid (\text{fst } e) \mid \dots \end{aligned}$$

Figure 3.2: Syntax of pretypes and preterms for a bidirectional elaborator.

These annotations are highly redundant, but it is far from clear how many of them can be mechanically reconstructed by our elaborator, nor if there is a consistent strategy for doing so. Users of typed functional programming languages like OCaml or Haskell might imagine that virtually all types can be inferred automatically; unfortunately, this is impossible in dependent type theory, for which type inference is undecidable [Dow93].

It turns out there is a fairly straightforward, local, and usable approach to type reconstruction known as *bidirectional type-checking* [Coq96; PT00; McB18; McB19]. The core insight of bidirectional type-checking is that for some preterms it is easy to reconstruct or *synthesize* its type (e.g., if we know a function’s type then we know the type of its applications), but for other preterms we must be given a type at which to *check* it (e.g., to type-check a function we need to be told the type of its input variable).

By explicitly splitting elaboration into two mutually-defined algorithms—type-checking and type synthesis—we can dramatically reduce type annotations. In fact, in Figure 3.2 we can see that our new preterm syntax has no type annotations whatsoever except for a single annotation form ($\text{chk } e \tau$) that we will use sparingly. The ebb and flow of information between terms and types—between checking and synthesis—leads to the eponymous bidirectional flow of information that has proven easily adaptable to new type theories. But when should we check, and when should we synthesize?

Slogan 3.2.7. *Types are checked in introduction rules, and synthesized in elimination rules.*

We replace our two algorithmic elaboration judgments $\Gamma \vdash \tau \text{ type} \rightsquigarrow A$ and $\Gamma \vdash e : A \rightsquigarrow a$ with three algorithmic judgments as follows:

1. $\Gamma \vdash \tau \Leftarrow \text{type} \rightsquigarrow A$ (“check τ ”) asserts that elaborating the pretype τ relative to $\vdash \Gamma \text{ cx}$ succeeds and produces the type $\Gamma \vdash A \text{ type}$.
2. $\Gamma \vdash e \Leftarrow A \rightsquigarrow a$ (“check e against A ”) asserts that elaborating the (unannotated) preterm e relative to $\vdash \Gamma \text{ cx}$ and a *given* type $\Gamma \vdash A \text{ type}$ succeeds with $\Gamma \vdash a : A$.
3. $\Gamma \vdash e \Rightarrow A \rightsquigarrow a$ (“synthesize A from e ”) asserts that elaborating the (unannotated) preterm e relative to $\vdash \Gamma \text{ cx}$ succeeds and *produces* both $\Gamma \vdash A \text{ type}$ and $\Gamma \vdash a : A$.

The first two judgments, $\Gamma \vdash \tau \Leftarrow \text{type} \rightsquigarrow A$ and $\Gamma \vdash e \Leftarrow A \rightsquigarrow a$, are similar to our previous judgments; when elaborating a preterm we are given a context and a type at which to check that preterm. In the third judgment, $\Gamma \vdash e \Rightarrow A \rightsquigarrow a$, we are also given

a preterm and a context, but we output *both a term and its type*. The arrows are meant to indicate the direction of information flow: when checking $e \Leftarrow A$ we are given A and using it to elaborate e , but when synthesizing $e \Rightarrow A$ we are producing A from e .

The rules for $\Gamma \vdash \tau \Leftarrow \text{type} \rightsquigarrow A$ are the same as those for $\Gamma \vdash \tau \text{ type} \rightsquigarrow A$, except that they reference the new checking judgment $\Gamma \vdash e \Leftarrow A \rightsquigarrow a$ instead of $\Gamma \vdash e : A \rightsquigarrow a$. But for each old $\Gamma \vdash e : A \rightsquigarrow a$ rule, we must decide whether this preterm should be checked or synthesized, and if the latter, how to reconstruct the type.

The easiest case is the variable ($\text{var } i$). Elaboration always takes place with respect to a context which records the types of each variable, so it is easy to synthesize the variable's type. Notably, unlike in our previous variable rule, we do not need to check type equality!

$$\frac{\Gamma = \Gamma'.A_i.A_{i-1} \cdots .A_0}{\Gamma \vdash (\text{var } i) \Rightarrow A_i[\mathbf{p}^{i+1}] \rightsquigarrow \mathbf{q}[\mathbf{p}^i]}$$

Next, let us consider the rules for Π -types. According to Slogan 3.2.7, the introduction form $(\text{lam } e)$ should be checked. As in Section 3.1, to check $\Gamma \vdash (\text{lam } e) \Leftarrow C \rightsquigarrow \lambda(b)$ we must recursively check the body of the lambda, $\Gamma.A \vdash e \Leftarrow B \rightsquigarrow b$. But where do A and B come from? (Last time, we elaborated them from lam 's annotations.) We might imagine that we can recover A and B from the given type C ,

$$\frac{\Gamma \vdash C \Leftrightarrow \Pi(A, B) \text{ type} \quad \Gamma.A \vdash e \Leftarrow B \rightsquigarrow b}{\Gamma \vdash (\text{lam } e) \Leftarrow C \rightsquigarrow \lambda(b)} \text{! ?}$$

but this rule does not make sense as written; $\Gamma \vdash C \Leftrightarrow D \text{ type}$ is an algorithm which takes two types and returns “yes” or “no”, and we cannot use it to invent the types A and B .

Worse yet, as foreshadowed in Remark 3.1.6, even if we can find A and B such that $\Gamma \vdash C \Leftrightarrow \Pi(A, B) \text{ type}$, there is no reason to expect this choice to be unique. That is, it could be that $\Gamma \vdash C \Leftrightarrow \Pi(A, B) \text{ type}$ and $\Gamma \vdash C \Leftrightarrow \Pi(A', B') \text{ type}$ both hold, but $A \neq A'$ (or alternatively, $A = A'$ and $B \neq B'$). If so, it is possible that e elaborates with respect to one of these choices but not the other, i.e., $\Gamma.A \vdash e \Leftarrow B \rightsquigarrow b$ succeeds but $\Gamma.A' \vdash e \Leftarrow B' \rightsquigarrow ?$ fails; even if both succeed, they will necessarily elaborate two different terms! We must foreclose these possibilities in order for elaboration to be well-defined.

Definition 3.2.8. A type theory has *injective Π -types* if $\Gamma \vdash \Pi(A, B) = \Pi(A', B') \text{ type}$ implies $\Gamma \vdash A = A' \text{ type}$ and $\Gamma.A \vdash B = B' \text{ type}$.

Definition 3.2.9. A type theory has *invertible Π -types* if it has injective Π -types and admits a computable function which, given $\Gamma \vdash C \text{ type}$, either produces the unique $\Gamma \vdash A \text{ type}$ and $\Gamma.A \vdash B \text{ type}$ for which $\Gamma \vdash C = \Pi(A, B) \text{ type}$, or determines that no such A, B exist.

Remark 3.2.10. That is, a type theory has injective Π -types if the type former $\Pi_\Gamma : (\sum_{A \in \text{Ty}(\Gamma)} \text{Ty}(\Gamma.A)) \rightarrow \text{Ty}(\Gamma)$ is injective. A type theory has invertible Π -types if the

image of Π_Γ is decidable and Π_Γ admits a (computable) partial inverse $\Pi_\Gamma^{-1} : \text{Im}(\Pi_\Gamma) \rightarrow (\sum_{A \in \text{Ty}(\Gamma)} \text{Ty}(\Gamma.A))$. \diamond

Particularly in light of Remark 3.2.10, one can easily extend the terminology of injectivity and invertibility to non- Π type formers.

Definition 3.2.11. If all the type constructors of a type theory are injective (resp., invertible), we say that the type theory *has injective (resp., invertible) type constructors*.

Having injective or invertible type constructors does not follow from normalization. (A type theory in which all empty types are equal may be normalizing but will not satisfy injectivity.) In practice, however, having invertible type constructors is almost always an immediate consequence of the *proof* of normalization. As we mentioned in Section 3.2.1, normalization proofs generally construct abstract syntax trees TyNf , TmNf of “ β -short, η -long” types and terms for which equality is both syntactic as well as sound and complete for judgmental equality. Given a type $\Gamma \vdash C$ type, we invert its head constructor by computing $\text{nfTy}(C) \in \text{TyNf}$, checking its head constructor in TyNf , and projecting its arguments.

Injectivity and invertibility are very strong conditions; function types in set theory are not injective, nor are Π -types injective in extensional type theory.

Exercise 3.5. Give an example of three sets X, Y, Z such that $X \not\cong Y$, but the set of functions $X \rightarrow Z$ is equal to the set of functions $Y \rightarrow Z$.

Exercise 3.6. We will see in Section 3.5 that type theory admits an interpretation in which closed types are sets. Exercise 3.5 shows that sets do not have injective Π -types, but these two facts together do not imply that type theory lacks injective Π -types. Why not?

Warning 3.2.12. In Section 3.5 we shall see that extensional type theory *does not* have injective type constructors, due to interactions between equality reflection and large elimination or universes (Theorem 3.5.19). We will proceed under the assumption that the theory we are elaborating has invertible type constructors, and in Chapter 4 we will see how to modify our type theory to substantiate this assumption.

Completing our elaborator The force of having invertible Π -types is to have an algorithm unPi which takes $\Gamma \vdash C$ type and returns the unique pair of types A, B for which $\Gamma \vdash C = \Pi(A, B)$ type, or raises an exception if this pair does not exist. Using unPi we can

repair our earlier attempt at checking $(\text{lam } e)$, and define the synthesis rule for $(\text{app } e_0 e_1)$:

$$\frac{\text{unPi}(C) = (A, B) \quad \Gamma.A \vdash e \Leftarrow B \rightsquigarrow b}{\Gamma \vdash (\text{lam } e) \Leftarrow C \rightsquigarrow \lambda(b)}$$

$$\frac{\Gamma \vdash e_0 \Rightarrow C \rightsquigarrow f \quad \text{unPi}(C) = (A, B) \quad \Gamma \vdash e_1 \Leftarrow A \rightsquigarrow a}{\Gamma \vdash (\text{app } e_0 e_1) \Rightarrow B[\text{id}.a] \rightsquigarrow \mathbf{app}(f, a)}$$

This is the only elaboration rule for $(\text{lam } e)$; in particular, there is no *synthesis* rule for lambda, because we cannot elaborate e without knowing what type A to add to the context. On the other hand, to synthesize the type of $(\text{app } e_0 e_1)$, we *synthesize* the type of e_0 ; if it is of the form $\Pi(A, B)$, we then *check* that e_1 has type A and then return B , suitably instantiated. Putting these rules together, the reader might notice that we cannot type-check $(\text{app } (\text{lam } e_0) e_1)$, because this would require *synthesizing* $(\text{lam } e_0)$. In fact, bidirectional type-checking cannot type-check β -redexes in general for this reason.

For this reason, we have included a *type-annotation* preterm $(\text{chk } e \tau)$ which allows users to explicitly annotate a preterm with a pretype. The type of this preterm is trivially synthesizable: it is the result of elaborating τ ! In order to synthesize $(\text{chk } e \tau)$, we simply *check* e against τ , and if successful, return that type.

$$\frac{\Gamma \vdash \tau \Leftarrow \text{type} \rightsquigarrow A \quad \Gamma \vdash e \Leftarrow A \rightsquigarrow a}{\Gamma \vdash (\text{chk } e \tau) \Rightarrow A \rightsquigarrow a}$$

In particular, we can type-check the β -redex from before, as long as we annotate the lambda with its intended type: $(\text{app } (\text{chk } (\text{lam } e_0) (\text{Pi } \tau_0 \tau_1)) e_1)$.

The above rule allows us to treat a checkable term as synthesizable. The converse is much easier: to *check* the type of a synthesizable term, we simply compare the synthesized type to the expected type.

$$\frac{\Gamma \vdash e \Rightarrow B \rightsquigarrow a \quad \Gamma \vdash A \Leftrightarrow B \text{ type}}{\Gamma \vdash e \Leftarrow A \rightsquigarrow a}$$

As written, the above rule applies to *any* checking problem because its conclusion is unconstrained. In our elaboration algorithm, we should only apply this rule if no other rule matches. It is the final “catch-all” clause for situations where we have not one but two sources of type information: on the one hand, we can synthesize e ’s type directly, and on the other hand, we are also given the type that e is supposed to have. Interestingly, this is the *only* rule where our bidirectional elaborator checks type equality $\Gamma \vdash A \Leftrightarrow B \text{ type}$.

Exercise 3.7. For each of $(\text{pair } e_0 e_1)$, $(\text{fst } e)$, and $(\text{snd } e)$, decide whether this preterm should be checked or synthesized, then write the algorithmic rule for elaborating it. (Hint: you must assume that Σ -types are invertible.)

3.3★ A case study in elaboration: definitions

To round out our discussion of elaboration, we sketch how to extend our concrete syntax and type-checker to account for *definitions*, a key part of any proof assistant. The input to a proof assistant is typically not a single term $e : \tau$ but a *sequence* of definitions

$$\begin{aligned} \text{def}_1 &: \tau_1 = e_1 \\ \text{def}_2 &: \tau_2 = e_2 \\ &\vdots \\ \text{def}_n &: \tau_n = e_n \end{aligned}$$

where every e_j and τ_j can mention def_i for $i < j$.

To account for this cross-definition dependency, we might imagine elaborating each definition one at a time, adding a new (nameless) variable to the context for each successful definition. Such a strategy might proceed as follows:

1. elaborate $\mathbf{1} \vdash \tau_1 \Leftarrow \text{type} \rightsquigarrow A_1$ and $\mathbf{1} \vdash e_1 \Leftarrow A_1 \rightsquigarrow a_1$; if successful,
2. elaborate $\mathbf{1}.A_1 \vdash \tau_2 \Leftarrow \text{type} \rightsquigarrow A_2$ and $\mathbf{1}.A_1 \vdash e_2 \Leftarrow A_2 \rightsquigarrow a_2$; if successful,
3. continue elaborating each τ_i and e_i in context $\mathbf{1}.A_1 \dots .A_{i-1}$ as above.

Unfortunately this algorithm is too naïve: if we treat def_1 as a *variable* of type A_1 , the type-checker will not have access to the definition $\text{def}_1 = a_1$. Consider:

$$\begin{aligned} \text{const} &: \text{Nat} \\ \text{const} &= 2 \\ \\ \text{proof} &: \text{const} \equiv 2 \\ \text{proof} &= \text{refl} \end{aligned}$$

Here `const` will successfully elaborate in the empty context to $\text{succ}(\text{succ}(\text{zero})) : \text{Nat}$, but the elaboration problem for `proof` will be $\mathbf{1}.\text{Nat} \vdash \text{refl} \Leftarrow \mathbf{q} \equiv \text{succ}(\text{succ}(\text{zero})) \rightsquigarrow ?$, which will fail: an arbitrary variable of type `Nat` is surely not equal to 2!

Remark 3.3.1. For readers familiar with functional programming, we summarize the above discussion as “**let** is no longer λ ,” in reference to the celebrated encoding of $(\text{let } x = a \text{ in } b)$ as $((\lambda x. b) a)$ often adopted in Lisp-family languages. This slogan is not unique to dependent type theory; users of ML-family languages may already be familiar with this phenomenon in light of the Hindley-Milner approach to typing **let**. \diamond

To solve this problem, we must somehow instrument our elaborator with the ability to remember not only the *type* of a definition but its *definiens* as well. There are several ways

to accomplish this; one possibility is to add a new form of *definitional context extension* “ $\Gamma.(q := a : A)$ ” in which the variable is judgmentally equal to a given term a [McB99; SP94]. We opt for an indirect but less invasive encoding of this idea: taking inspiration from Section 2.6.2, wherein we encoded “extending the context by a type variable” by adding a new type U whose terms are codes for types, we will add a new type former, *singleton types*, whose terms are elements of A judgmentally equal to a .

Singleton types The singleton type of $\Gamma \vdash a : A$, written $\Gamma \vdash \mathbf{Sing}(A, a)$ type, is a type whose elements are in bijection with the elements of $\mathbf{Tm}(\Gamma, A)$ that are equal to a , namely the singleton subset $\{a\}$ [Asp95; SH06]. That is, naturally in Γ ,

$$\begin{aligned} \mathbf{Sing}_\Gamma &: (\sum_{A \in \mathbf{Ty}(\Gamma)} \mathbf{Tm}(\Gamma, A)) \rightarrow \mathbf{Ty}(\Gamma) \\ \iota_{\Gamma, A, a} &: \mathbf{Tm}(\Gamma, \mathbf{Sing}(A, a)) \cong \{b \in \mathbf{Tm}(\Gamma, A) \mid b = a\} \end{aligned}$$

In inference rules,

$$\begin{array}{c} \frac{\Gamma \vdash a : A}{\Gamma \vdash \mathbf{Sing}(A, a) \text{ type}} \qquad \frac{\Gamma \vdash a : A}{\Gamma \vdash \mathbf{in}(a) : \mathbf{Sing}(A, a)} \qquad \frac{\Gamma \vdash s : \mathbf{Sing}(A, a)}{\Gamma \vdash \mathbf{out}(s) : A} \\ \\ \frac{\Gamma \vdash s : \mathbf{Sing}(A, a)}{\Gamma \vdash \mathbf{out}(s) = a : A} \qquad \frac{\Gamma \vdash a : A \quad \Gamma \vdash s : \mathbf{Sing}(A, a)}{\Gamma \vdash \mathbf{in}(\mathbf{out}(s)) = s : \mathbf{Sing}(A, a)} \end{array}$$

This definition may seem rather odd, but note that a variable of type $\mathbf{Sing}(A, a)$ determines a term $\mathbf{out}(q) : A[\mathbf{p}]$ that is judgmentally equal to $a[\mathbf{p}]$, thereby allowing us to extend contexts by “defined variables.”

Remark 3.3.2. In extensional type theory, we can define singleton types as pairs of an element of A and a proof that this element equals a , i.e., $\mathbf{Sing}(A, a) := \Sigma(A, \mathbf{Eq}(A[\mathbf{p}], q, a[\mathbf{p}]))$ with $\mathbf{in}(a) := \mathbf{pair}(a, \mathbf{refl})$ and $\mathbf{out}(s) := \mathbf{fst}(s)$. This encoding makes essential use of equality reflection, but singleton types can also be added as a primitive type former to type theories without equality reflection, without disrupting normalization. \diamond

Extending our elaborator We begin by introducing concrete syntax for lists of $e : \tau$ pairs, which we call *declarations*:

$$\begin{array}{ll} \text{Declarations} & ds := (\text{decls } (e_1 \tau_1) \dots) \\ \text{Pretypes} & \tau := \dots \\ \text{Preterms} & e := \dots \end{array}$$

We extend our bidirectional elaborator as follows. First, we parameterize all our judgments by a second context Θ that keeps track of which variables in Γ are ordinary

“local” variables (introduced by types/terms such as Π or λ), and which variables refer to declarations. We write Θ as a list $1.\text{decl}.\text{decl}.\text{local}.\dots$ with the same length as $\Gamma = 1.A_1.A_2.A_3.\dots$, to indicate in this case that only the variable of type A_3 is local. We will replace the variable rule shortly; the remaining elaboration rules do not interact with Θ except to extend Θ by local whenever a new variable is added to the context Γ .

Secondly, we introduce a new algorithmic judgment $\Gamma; \Theta \vdash ds \text{ ok}$ which type-checks a list of declarations ds by elaborating the first declaration ($e_1 \tau_1$) in context $\Gamma; \Theta$ into the term $a_1 : A_1$, and then elaborating the remaining declarations in context $\Gamma.\text{Sing}(A_1, a_1); \Theta.\text{decl}$.

$$\frac{}{\Gamma; \Theta \vdash (\text{decls}) \text{ ok}} \quad \frac{\Gamma; \Theta \vdash \tau_1 \Leftarrow \text{type} \rightsquigarrow A_1 \quad \Gamma; \Theta \vdash e_1 \Leftarrow A_1 \rightsquigarrow a_1 \quad \Gamma.\text{Sing}(A_1, a_1); \Theta.\text{decl} \vdash (\text{decls } (e_2 \tau_2) \dots) \text{ ok}}{\Gamma; \Theta \vdash (\text{decls } (e_1 \tau_1) (e_2 \tau_2) \dots) \text{ ok}}$$

Finally, we must edit our variable rule to account for whether a variable is an ordinary local variable or refers to an earlier declaration; in the latter case, we must insert an extra $\text{out}(-)$ around the variable so it has the correct type A rather than $\text{Sing}(A, a)$.

$$\frac{\Gamma = \Gamma'.A_i.A_{i-1}.\dots.A_0 \quad \Theta = \Theta'.\text{local}.x_{i-1}.\dots.x_0}{\Gamma; \Theta \vdash (\text{var } i) \Rightarrow A_i[\mathbf{p}^{i+1}] \rightsquigarrow \mathbf{q}[\mathbf{p}^i]} \quad \frac{\Gamma = \Gamma'.A_i.A_{i-1}.\dots.A_0 \quad \Theta = \Theta'.\text{decl}.x_{i-1}.\dots.x_0 \quad \text{unSing}(A_i) = (A, a)}{\Gamma; \Theta \vdash (\text{var } i) \Rightarrow A[\mathbf{p}^{i+1}] \rightsquigarrow \text{out}(\mathbf{q}[\mathbf{p}^i])}$$

In the second rule above, the rules of singleton types ensure that the elaborated term $\text{out}(\mathbf{q}[\mathbf{p}^i])$ is judgmentally equal to $a[\mathbf{p}^{i+1}]$, where a is the previously-elaborated definiens of the corresponding declaration. Putting everything together, to check an input file $(\text{decls } (e_1 \tau_1) (e_2 \tau_2) \dots)$ we attempt to derive $1; 1 \vdash (\text{decls } (e_1 \tau_1) (e_2 \tau_2) \dots) \text{ ok}$.

To sum up, we emphasize once again that although this book focuses on the *core calculi* of proof assistants, it is impossible to have a satisfactory understanding of this topic without paying heed to their *surface languages* as well; often, the best way to understand a new surface language feature is to add a new feature in the core language to accommodate it. Ideally, our alterations to the core language will be minor but will significantly simplify elaboration.

3.4 Metatheory for computing

Our focus on type-checking has led us to normalization (Definition 3.2.4) and invertible type constructors (Definition 3.2.11) as metatheorems essential to the implementation of type theory. Notably, these metatheorems are stated with respect to types and terms in *arbitrary* contexts; in this section, we will discuss two more important metatheorems that

concern only terms in the empty context $\mathbf{1}$, namely *consistency* and *canonicity*. Neither of these properties is needed to implement a type-checker, but as we will see, they are essential to the applications of type theory to logic and programming languages respectively.

Definition 3.4.1. A type theory is *consistent* if there is no closed term $\mathbf{1} \vdash a : \mathbf{Void}$.

Consistency is the lowest bar that a type theory must pass in order to function as a logic. When we interpret types as logical propositions, \mathbf{Void} corresponds to the false proposition. By the rules of \mathbf{Void} (Section 2.5.1), the existence of a closed term $\mathbf{1} \vdash a : \mathbf{Void}$ (an assumption-free proof of false) implies that every closed type has at least one closed term $\mathbf{1} \vdash \mathbf{absurd}(a) : A$, or in other words, that every proposition has a proof. Thus Definition 3.4.1 corresponds to logical consistency in the traditional sense.

At this point we pause to sketch the model theory of type theory. In Chapter 2 we were careful to formulate the judgments of type theory as (indexed) sets, and the rules of type theory as (dependently-typed) operations between these sets and equations between these operations. As a result we can regard this data as a kind of generalized *algebra signature*, in the sense of Section 2.5.3; in particular, we obtain a general notion of “implementation” of, or *algebra* for, this signature—more commonly known as a *model of type theory*.

Definition 3.4.2. A *model of type theory* \mathcal{M} consists of the following data:

1. a set $\mathbf{Cx}_{\mathcal{M}}$ of \mathcal{M} -contexts,
2. for each $\Delta, \Gamma \in \mathbf{Cx}_{\mathcal{M}}$, a set $\mathbf{Sb}_{\mathcal{M}}(\Delta, \Gamma)$ of \mathcal{M} -substitutions from Δ to Γ ,
3. for each $\Gamma \in \mathbf{Cx}_{\mathcal{M}}$, a set $\mathbf{Ty}_{\mathcal{M}}(\Gamma)$ of \mathcal{M} -types in Γ , and
4. for each $\Gamma \in \mathbf{Cx}_{\mathcal{M}}$ and $A \in \mathbf{Ty}_{\mathcal{M}}(\Gamma)$, a set $\mathbf{Tm}_{\mathcal{M}}(\Gamma, A)$ of \mathcal{M} -terms of A in Γ ,
5. an *empty* \mathcal{M} -context $\mathbf{1}_{\mathcal{M}} \in \mathbf{Cx}_{\mathcal{M}}$,
6. for each $\Gamma \in \mathbf{Cx}_{\mathcal{M}}$ and $A \in \mathbf{Ty}_{\mathcal{M}}(\Gamma)$, an \mathcal{M} -context extension $\Gamma.\mathcal{M}A \in \mathbf{Cx}_{\mathcal{M}}$,
7. for $\Gamma \in \mathbf{Cx}_{\mathcal{M}}$, $A \in \mathbf{Ty}_{\mathcal{M}}(\Gamma)$, and $B \in \mathbf{Ty}_{\mathcal{M}}(\Gamma.\mathcal{M}A)$, an \mathcal{M} - Π type $\mathbf{\Pi}_{\mathcal{M}}(A, B) \in \mathbf{Ty}_{\mathcal{M}}(\Gamma)$,
8. and every other context, substitution, type, and term forming operation described in Appendix A, all subject to all the equations stated in Appendix A.

Definition 3.4.3. Given two models of type theory \mathcal{M}, \mathcal{N} , a *homomorphism of models of type theory* $f : \mathcal{M} \rightarrow \mathcal{N}$ consists of the following data:

1. a function $\mathbf{Cx}_f : \mathbf{Cx}_{\mathcal{M}} \rightarrow \mathbf{Cx}_{\mathcal{N}}$,
2. for each $\Delta, \Gamma \in \mathbf{Cx}_{\mathcal{M}}$, a function $\mathbf{Sb}_f(\Delta, \Gamma) : \mathbf{Sb}_{\mathcal{M}}(\Delta, \Gamma) \rightarrow \mathbf{Sb}_{\mathcal{N}}(\mathbf{Cx}_f(\Delta), \mathbf{Cx}_f(\Gamma))$,

3. for each $\Gamma \in \text{Cx}_{\mathcal{M}}$, a function $\text{Ty}_f(\Gamma) : \text{Ty}_{\mathcal{M}}(\Gamma) \rightarrow \text{Ty}_{\mathcal{N}}(\text{Cx}_f(\Gamma))$, and
4. for each $\Gamma \in \text{Cx}_{\mathcal{M}}$ and $A \in \text{Ty}_{\mathcal{M}}(\Gamma)$, a function $\text{Tm}_f(\Gamma, A) : \text{Tm}_{\mathcal{M}}(\Gamma, A) \rightarrow \text{Tm}_{\mathcal{N}}(\text{Cx}_f(\Gamma), \text{Ty}_f(\Gamma)(A))$,
5. such that $\text{Cx}_f(\mathbf{1}_{\mathcal{M}}) = \mathbf{1}_{\mathcal{N}}$,
6. and every other context, substitution, type, and term forming operation of \mathcal{M} is also sent to the corresponding operation of \mathcal{N} in a similar fashion.

Definition 3.4.4. The sets Cx , $\text{Sb}(\Delta, \Gamma)$, $\text{Ty}(\Gamma)$, and $\text{Tm}(\Gamma, A)$, equipped with the context, substitution, type, and term forming operations described in Appendix A, tautologically form a model of type theory \mathcal{T} known as the *syntactic model*.

Theorem 3.4.5. *The syntactic model \mathcal{T} is the initial model of type theory; that is, for any model of type theory \mathcal{M} , there exists a unique homomorphism of models $\mathcal{T} \rightarrow \mathcal{M}$.*

The notions of model and homomorphism are quite complex, but they are mechanically derivable from the rules of type theory as presented in Appendix A, viewed as the signature of a quotient inductive-inductive type [KKA19] or generalized algebraic theory [Car86]. The initiality of the syntactic model expresses the fact that type theory is the “least” model of type theory, in the sense that it—by definition—satisfies all the rules of type theory and no others; this mirrors the sense in which initiality of \mathbb{N} with respect to $(1 + -)$ -algebras expresses that the natural numbers are generated by **zero** and **suc**($-$).

Remark 3.4.6. Theorem 3.4.5 should be regarded as stating the soundness and completeness of type theory with respect to this notion of model. The homomorphism $\mathcal{T} \rightarrow \mathcal{M}$ expresses soundness: the syntax of type theory can be interpreted into any model \mathcal{M} . Conversely, the fact that the syntax constitutes a model \mathcal{T} expresses completeness: any result that holds for all models must in particular hold for \mathcal{T} and thence for the syntax.

We note that Definitions 3.4.2 and 3.4.3 were carefully chosen so as to make soundness and completeness nearly tautological, and indeed, this is evidenced by the fact that these definitions and theorems can be mechanically derived by the general machinery of quotient inductive-inductive types or generalized algebraic theories. Unimpressed readers may commiserate with Girard’s “broccoli logic” critique of such semantics [Gir99]. \diamond

While the definition of a model does not lend much insight into type theory on its own, the model theory of type theory is an essential tool in the metatheorist’s toolbox; to prove any property of the syntactic model \mathcal{T} , we simply produce a model of type theory \mathcal{M} such that Theorem 3.4.5 implies the property in question. In the case of consistency, it suffices to exhibit any non-trivial model of type theory whatsoever.

Theorem 3.4.7. *Suppose there exists a model of type theory \mathcal{M} such that $\text{Tm}_{\mathcal{M}}(\mathbf{1}_{\mathcal{M}}, \mathbf{Void}_{\mathcal{M}})$ is empty; then type theory is consistent.*

Proof. We must show that from the existence of \mathcal{M} and a term $a \in \text{Tm}(\mathbf{1}, \mathbf{Void})$ we can derive a contradiction. By Theorem 3.4.5, there is a homomorphism of models $f : \mathcal{T} \rightarrow \mathcal{M}$, and in particular a function $\text{Tm}_f(\mathbf{1}, \mathbf{Void}) : \text{Tm}(\mathbf{1}, \mathbf{Void}) \rightarrow \text{Tm}_{\mathcal{M}}(\mathbf{1}_{\mathcal{M}}, \mathbf{Void}_{\mathcal{M}})$; applying this function to a produces an element of $\text{Tm}_{\mathcal{M}}(\mathbf{1}_{\mathcal{M}}, \mathbf{Void}_{\mathcal{M}})$, an empty set. \square

In Section 3.5 we will see that there is a “standard” set-theoretic model \mathcal{S} of extensional type theory in which contexts are sets, types are families of sets indexed by their context, and each type former is interpreted as the corresponding construction on indexed sets. As a trivial corollary of this model and Theorem 3.4.7, we obtain the consistency of extensional type theory. We postpone further details of the set-theoretic model to Section 3.5; interested readers may also consult Castellan, Clairambault, and Dybjer [CCD21] and Hofmann [Hof97] for tutorials on the categorical semantics of type theory.

Theorem 3.4.8 (Martin-Löf [ML84b]). *Extensional type theory is consistent.*

Note that while an inconsistent type theory is useless as a logic, it may still be useful for programming; indeed, many modern functional programming languages include some limited forms of dependent types despite being inconsistent.

Exercise 3.8. Consider an unrestricted fixed-point operator $\text{fix} : (A \rightarrow A) \rightarrow A$, i.e.,

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash a : A[\mathbf{p}]}{\Gamma \vdash \text{fix}(a) : A} \quad \text{✎}$$

Show that adding such a rule results in an inconsistent type theory.

In fact, our final metatheorem is directly connected to the interpretation of type theory as a programming language, although the connection may not be immediately apparent.

Definition 3.4.9. A type theory enjoys *canonicity* if for every closed $\mathbf{1} \vdash b : \mathbf{Bool}$ either $\mathbf{1} \vdash b = \mathbf{true} : \mathbf{Bool}$ or $\mathbf{1} \vdash b = \mathbf{false} : \mathbf{Bool}$, but not both.

Remark 3.4.10. Another common statement of canonicity is that for every closed $\mathbf{1} \vdash n : \mathbf{Nat}$ either $\mathbf{1} \vdash n = \mathbf{zero} : \mathbf{Nat}$ or $\mathbf{1} \vdash n = \mathbf{suc}(m) : \mathbf{Nat}$ where $\mathbf{1} \vdash m : \mathbf{Nat}$. This statement is not equivalent to Definition 3.4.9 in general, but in practice one only considers type theories that satisfy both or neither, and proofs of one also imply the other *en passant*. \diamond

Remark 3.4.11. Consistency states that $\text{Tm}(\mathbf{1}, \mathbf{Void}) \cong \emptyset$, whereas canonicity states that $\text{Tm}(\mathbf{1}, \mathbf{Bool}) \cong \{\star, \star'\}$ and $\text{Tm}(\mathbf{1}, \mathbf{Nat}) \cong \mathbb{N}$. As discussed at length in Section 2.5, none of these properties hold in Γ because variables can produce noncanonical terms at any type; however, there are indeed no noncanonical *closed* terms of type \mathbf{Void} , \mathbf{Bool} , or \mathbf{Nat} . \diamond

Theorem 3.4.12. *Extensional type theory enjoys canonicity.*

Frustratingly, although Theorem 3.4.12 was certainly known to researchers in the 1970s and 1980s, the authors are unable to locate a precise reference from this time period. For a modern proof of Theorem 3.4.12, we refer the reader to Sterling [Ste21, Chapter 4].

Like consistency—and normalization and invertibility of type constructors—canonicity can be established by constructing a model of type theory, although the proofs of the latter three metatheorems are considerably more involved than the proof of consistency. Canonicity models interpret the contexts, substitutions, types, and terms of type theory as pairs of that syntactic object along with additional data which explains how that object may be placed in canonical form [Fre78; LS88; MS93; Cro94; Fio02; AK16; Coq19; KHS19]. Such models can be seen as *displayed models* of type theory over the syntactic model, and are called *gluing models* in the categorical literature. The interested reader may consult Lambek and Scott [LS88] for information on this perspective as it applies to higher-order logic, and Crole [Cro94] for an application to simple type theory.

Exercise 3.9. In light of Remark 3.4.11, we might imagine that canonicity follows from the existence of a model of type theory \mathcal{M} for which $\text{Tm}_{\mathcal{M}}(\mathbf{1}_{\mathcal{M}}, \mathbf{Bool}_{\mathcal{M}})$ has exactly two elements. This is not the case; why? (Why can't we mimic the proof of Theorem 3.4.7?)

The force of canonicity is that it implies the existence of an “evaluation” algorithm that, given a closed boolean $\mathbf{1} \vdash a : \mathbf{Bool}$, reports whether a is equal to **true** or to **false**. There are two ways to obtain such an algorithm; the first is to prove canonicity in a constructive metatheory, so that the proof itself constitutes such an algorithm. The second is to appeal to Markov's principle: because derivation trees are recursively enumerable, a classical proof of canonicity implies that the naïve enumeration algorithm will terminate.

In a direct sense, such an algorithm is indeed an *interpreter* for closed terms of type theory. But canonicity also produces a much richer notion of computational adequacy for type theory; giving this theory its due weight would take us too far afield, but we will briefly sketch the highlights. By results in categorical realizability [Jac99; Oos08], essentially every model of computation gives rise to a highly structured and well-behaved category known as a *realizability topos*; these categories support models of dependent type theory in which terms of type **Bool** are (equivalence classes of) boolean computations in some idealized model of computation. For instance, in the *effective topos* [Hyl82], closed terms of type **Bool** are equivalence classes of Turing machines modulo Kleene equivalence (i.e., two machines are equivalent if they coterminate with the same value).

Because models of type theory in realizability topoi interpret terms in concrete (albeit theoretical) notions of computation such as Turing machines or combinator calculi, they can be regarded abstractly as *compilers* for type theory. Alternatively, they serve to justify the *program extraction* mechanisms found in proof assistants such as Coq and Agda, which associate to each term an OCaml or Haskell program whose observable behavior is compatible with the definitional equality of type theory.

From this perspective, canonicity guarantees that definitional equality fully constrains the observable behaviors of extracted programs: for any closed boolean $1 \vdash b : \mathbf{Bool}$, every possible extract for b must evaluate to (the extract of) either **true** or **false**, as predetermined by whether $b = \mathbf{true}$ or $b = \mathbf{false}$. Note that it is still possible for two different extracts of b to have very different execution traces; canonicity only constrains their observable behavior, considered modulo some appropriate notion of observational equivalence.

Remark 3.4.13. The above discussion may clarify why canonicity is harder to prove than consistency: consistency implies the *existence* of a non-trivial model of type theory, whereas canonicity places a constraint on *all* models of type theory. \diamond

We emphasize once more that, unlike normalization and invertibility of type constructors, neither consistency nor canonicity is required to implement a bidirectional type-checker for type theory. However, it seems safe to assume that anybody writing such a type-checker is interested in type theory’s applications to logic or programming or both, in which case consistency and canonicity are relevant properties. In addition, failures of canonicity often indicate a paucity of definitional equalities that can have a negative effect on the usability of a type theory even as a logic.

3.5★ *The set model of type theory*

We now spell out the details of the set-theoretic model \mathcal{S} of extensional type theory alluded to in Section 3.4 [Hof97]. The remainder of this book will not depend on this section, but it may nevertheless be valuable to readers interested in better understanding the model theory of type theory or how type theory relates to traditional mathematics.

In short, \mathcal{S} interprets the contexts of type theory as sets, substitutions as functions, dependent types as indexed families of sets, terms as indexed families of elements, and every type- and term-forming operation as its “standard” mathematical counterpart. For example, the \mathcal{S} -interpretation of the closed functions from \mathbf{Nat} to \mathbf{Nat} , $\mathbf{Tm}_{\mathcal{S}}(\mathbf{1}_{\mathcal{S}}, \Pi_{\mathcal{S}}(\mathbf{Nat}_{\mathcal{S}}, \mathbf{Nat}_{\mathcal{S}}))$, is (isomorphic to) the set of ordinary mathematical functions $\mathbb{N} \rightarrow \mathbb{N}$.

The main subtlety in defining \mathcal{S} is that we would like the set $\mathbf{Cx}_{\mathcal{S}}$ of \mathcal{S} -contexts to be “the collection of all sets,” but this collection is unfortunately not a set: by Russell’s paradox, having a “set of all sets including itself” leads to contradiction. To properly circumvent this issue we must introduce the notion of *Grothendieck universes*, the set-theoretic cousins of the type-theoretic universes introduced in Section 2.6.

3.5.1 *Grothendieck universes*

Grothendieck universes are sets that resemble a “set of all sets” without falling victim to Russell’s paradox. Roughly speaking, they are collections of sets that are closed under all

the operations of set theory: they contain \emptyset and are closed under formation of powersets, unions, set comprehensions, and so forth.

Definition 3.5.1. A *Grothendieck universe* \mathcal{V} is a set satisfying the following conditions:

1. $\emptyset \in \mathcal{V}$.
2. *Transitivity:* If $X \in \mathcal{V}$ and $Y \in X$, then $Y \in \mathcal{V}$.
3. *Closure under powersets:* If $X \in \mathcal{V}$ then $\mathcal{P}(X) \in \mathcal{V}$.
4. *Closure under indexed unions:* If $X \in \mathcal{V}$ and $f : X \rightarrow \mathcal{V}$, then $\bigcup_{x \in X} f(x) \in \mathcal{V}$.
5. $\mathbb{N} \in \mathcal{V}$. (This condition is omitted by many authors.)

We admit that Definition 3.5.1 may seem somewhat mysterious; unfortunately, thoroughly justifying these axioms is beyond the scope of this book. We refer the reader to Shulman [Shu08] for a reference which assumes relatively little set-theoretic background.

For our purposes, the axioms of Grothendieck universes satisfy three important properties. First, all the closure properties of Grothendieck universes are closure properties of sets: replacing $X \in \mathcal{V}$ with “ X is a set,” it is true that \emptyset and \mathbb{N} are sets, and that sets are transitive and closed under powersets and indexed unions. In other words, the collection of all sets looks like a Grothendieck universe—except that a Grothendieck universe must be a set, which the collection of all sets is not.

Secondly, these closure conditions imply all the other usual closure conditions of sets. For example, \mathcal{V} is also closed under subsets, products, and function spaces, defined by their standard set-theoretic encodings. We prove a number of these closure conditions below, noting that these are not intended to be exhaustive.

Lemma 3.5.2. *Every Grothendieck universe \mathcal{V} is closed under the following constructions:*

1. Subsets: If $X \in \mathcal{V}$ and $Y \subseteq X$, then $Y \in \mathcal{V}$.
2. Binary unions: If $X, Y \in \mathcal{V}$ then $X \cup Y \in \mathcal{V}$.
3. Products: If $X, Y \in \mathcal{V}$ then $X \times Y \in \mathcal{V}$.
4. Function spaces: If $X, Y \in \mathcal{V}$ then $X \rightarrow Y \in \mathcal{V}$.
5. Indexed coproducts: If $X \in \mathcal{V}$ and $f : X \rightarrow \mathcal{V}$, then $\sum_{x \in X} f(x) \in \mathcal{V}$.
6. Indexed products: If $X \in \mathcal{V}$ and $f : X \rightarrow \mathcal{V}$, then $\prod_{x \in X} f(x) \in \mathcal{V}$.

Proof.

1. This follows directly from $Y \in \mathcal{P}(X) \in \mathcal{V}$ and transitivity.
2. We obtain binary unions as a special case of indexed unions, using the fact that the two-element set $\mathcal{P}(\mathcal{P}(\emptyset)) = \{\emptyset, \{\emptyset\}\}$ is an element of \mathcal{V} . Let $f : \mathcal{P}(\mathcal{P}(\emptyset)) \rightarrow \mathcal{V}$ be the function sending \emptyset to X and $\{\emptyset\}$ to Y ; then we define $X \cup Y := \bigcup_{x \in \mathcal{P}(\mathcal{P}(\emptyset))} f(x) \in \mathcal{V}$.
3. Following the usual set-theoretic construction, we define $X \times Y$ to be the subset of $\mathcal{P}(\mathcal{P}(X \cup Y))$ consisting of *ordered pairs* (x, y) with $x \in X$ and $y \in Y$, where $(x, y) := \{\{x\}, \{x, y\}\}$. We observe that $X \times Y \in \mathcal{V}$ by the closure of \mathcal{V} under binary unions, powersets, and subsets.
4. Functions $f : X \rightarrow Y$ are in bijection with subsets $S \subseteq X \times Y$ satisfying the condition that for all $x \in X$, there exists a unique $y \in Y$ such that the ordered pair (x, y) is in S . We may therefore take the collection of all such S —a subset of $\mathcal{P}(X \times Y)$ and thus an element of \mathcal{V} —as the definition of the function space $X \rightarrow Y$.
5. We define the indexed disjoint union $\sum_{x \in X} f(x)$ as the subset of $X \times \bigcup_{x \in X} f(x)$ consisting of ordered pairs (x, y) for which $y \in f(x)$.
6. Similarly, we define the indexed product $\prod_{x \in X} f(x)$ as the subset of $X \rightarrow \bigcup_{x \in X} f(x)$ consisting of the functions g for which $g(x) \in f(x)$ for all $x \in X$. \square

Finally and most importantly, although the existence of Grothendieck universes is independent from the axioms of ordinary (ZFC) set theory, it is consistent to assume that they exist,³ and the resulting theory is well-understood albeit stronger than ZFC.

Advanced Remark 3.5.3. In fact, assuming the existence of a Grothendieck universe \mathcal{V} is exactly the same as assuming the existence of a strongly inaccessible cardinal. This is fairly modest as far as large cardinal axioms are concerned, but it is strong enough that ZFC+ \mathcal{V} proves Con(ZFC). Indeed, \mathcal{V} is a model of ZFC! \diamond

Remark 3.5.4. As we will see in Section 3.5.4, one consequence of the set-theoretic model of type theory is the consistency of type theory. By Gödel’s incompleteness theorem, constructing this model must require a metatheory stronger than extensional type theory. Although ZFC and extensional type theory are not exactly aligned in strength, we should not be surprised that plain ZFC is too weak. In fact, if we augment extensional type theory with an impredicative universe of propositions (Section 2.8) and a few axioms, it becomes exactly as strong as ZFC with a universe hierarchy [Wer97]. \diamond

In the remainder of Section 3.5, we will rely on an ambient assumption that there is a $(\omega + 1)$ -indexed hierarchy of nested Grothendieck universes, in the following sense.

³In particular, it does not follow from the axioms that \mathcal{V} contains itself.

Definition 3.5.5. For a partial order I , an I -hierarchy of Grothendieck universes $(\mathcal{V}_i)_{i \in I}$ is a family of Grothendieck universes \mathcal{V}_i such that $\mathcal{V}_i \in \mathcal{V}_j$ whenever $i < j$.

Axiom 3.5.6. There exists an $(\omega + 1)$ -hierarchy of Grothendieck universes $\mathcal{V}_0 \in \dots \in \mathcal{V}_\omega$.

Intuitively, Axiom 3.5.6 states that \mathcal{V}_0 contains all the sets that exist in ZFC, \mathcal{V}_1 contains all the sets of ZFC+ \mathcal{V}_0 , \mathcal{V}_2 contains all the sets of ZFC+ \mathcal{V}_0 + \mathcal{V}_1 , and so forth. One often refers to the sets of ZFC as *small sets* for emphasis, and in general for a Grothendieck universe \mathcal{V} we say that a set X is \mathcal{V} -small if $X \in \mathcal{V}$. Thus Axiom 3.5.6 equivalently states that small sets are \mathcal{V}_i -small and \mathcal{V}_i is \mathcal{V}_j -small for all $i < j$.

3.5.2 The substitution calculus of sets

Exhibiting a model \mathcal{M} of type theory (Definition 3.4.2) requires an enormous amount of data, but we can break the process down into three steps:

1. First, one must define the sets of \mathcal{M} -contexts $Cx_{\mathcal{M}}$, \mathcal{M} -substitutions $Sb_{\mathcal{M}}(-, -)$, \mathcal{M} -types $Ty_{\mathcal{M}}(-)$, and \mathcal{M} -terms $Tm_{\mathcal{M}}(-, -)$.
2. Next, one must provide the \mathcal{M} -interpretations of the rules of the substitution calculus (Section 2.3), the core structure of type theory governing variables and substitutions, and verify that these satisfy the associated equations.
3. Finally, for each connective (Π -types, **Void**, U_i , etc.) one provides \mathcal{M} -interpretations of the associated rules, and again verifies the associated equations.

The steps must be performed in this order, because the choice of sets (e.g., $Cx_{\mathcal{M}}$) in the first step affects the interpretation of the substitution calculus (e.g., $p_{\mathcal{M}}$) in the second step, which in turn affects the interpretation of every connective. However, the interpretations of non- U connectives do not depend on one another and can be added in any order, because we were careful in Chapter 2 to avoid mentioning (e.g.) Π -types in the rules for Σ -types.

We will now carry out the first two steps of defining the set model \mathcal{S} . By the end of this subsection, we will have a model of a dependent type theory with no connectives, mirroring the situation at the end of Section 2.3.

The basic sets With the machinery of Grothendieck universes (Definition 3.5.1) under our belt, we can now define the basic sets of the \mathcal{S} -interpretation of type theory: the \mathcal{S} -contexts, \mathcal{S} -substitutions, \mathcal{S} -types, and \mathcal{S} -terms. Rather than defining the set of \mathcal{S} -contexts $Cx_{\mathcal{S}}$ to be the nonexistent “set of all sets,” we will define it to be a Grothendieck universe, a set of *some* sets which is closed under all the set-forming operations of set

theory. For reasons that will become clear later, we choose the set of \mathcal{S} -contexts to be \mathcal{V}_ω , the largest Grothendieck universe asserted by Axiom 3.5.6.

$$\mathbf{Cx}_{\mathcal{S}} := \mathcal{V}_\omega$$

For any two \mathcal{S} -contexts $\Delta, \Gamma \in \mathbf{Cx}_{\mathcal{S}}$, the set of \mathcal{S} -substitutions from Δ to Γ is simply the set of ordinary functions from Δ to Γ :

$$\mathbf{Sb}_{\mathcal{S}}(\Delta, \Gamma) := \Delta \rightarrow \Gamma \quad (\Delta, \Gamma \in \mathbf{Cx}_{\mathcal{S}})$$

Notation 3.5.7. Throughout this section, the variables $\Gamma, \gamma, A, a, \dots$ range over \mathcal{S} -contexts, substitutions, types, and terms, *not* syntactic contexts, substitutions, types, and terms as they generally have throughout this book. We believe this notation is the least confusing in the long run, but the reader should proceed cautiously.

Intuitively, an \mathcal{S} -type A in \mathcal{S} -context Γ should be a family of sets indexed by the set Γ , *i.e.*, a choice of set $A(x)$ for each $x \in \Gamma$. As in our definition of $\mathbf{Cx}_{\mathcal{S}}$, we can obtain a set of such families by restricting all the sets $A(x)$ to be elements of \mathcal{V}_ω :

$$\mathbf{T}_{\mathcal{S}}(\Gamma) := \Gamma \rightarrow \mathcal{V}_\omega \quad (\Gamma \in \mathbf{Cx}_{\mathcal{S}})$$

Finally, given an \mathcal{S} -context $\Gamma \in \mathcal{V}_\omega$ and an \mathcal{S} -type $A : \Gamma \rightarrow \mathcal{V}_\omega$ in that context, an \mathcal{S} -term $a \in \mathbf{Tm}_{\mathcal{S}}(\Gamma, A)$ should be a family of elements of each $A(x)$ for each $x \in \Gamma$. In other words, a should be a dependent function $(x : \Gamma) \rightarrow A(x)$, where $a(x) \in A(x)$ for all $x \in \Gamma$. Set-theoretically, such functions are more commonly understood as elements of the Γ -indexed product of the sets $A(-)$; see Remarks 2.4.1 and 2.4.5.

$$\mathbf{Tm}_{\mathcal{S}}(\Gamma, A) := \prod_{x \in \Gamma} A(x) \quad (\Gamma \in \mathbf{Cx}_{\mathcal{S}}, A \in \mathbf{T}_{\mathcal{S}}(\Gamma))$$

Summing up, we define \mathcal{S} -contexts as (\mathcal{V}_ω -small) sets, \mathcal{S} -substitutions as functions, \mathcal{S} -types as indexed families of (\mathcal{V}_ω -small) sets, and \mathcal{S} -terms as indexed families of elements.

The category of substitutions Having now defined the basic sets underlying the \mathcal{S} -interpretation of type theory, our next task is to define the operations of the substitution calculus (collected in the first section of Appendix A), starting with the identity and composition of substitutions.

For every \mathcal{S} -context $\Gamma \in \mathbf{Cx}_{\mathcal{S}}$, we must define an identity \mathcal{S} -substitution $\mathbf{id}_{\mathcal{S}}$ in $\mathbf{Sb}_{\mathcal{S}}(\Gamma, \Gamma)$. Unfolding the definitions of $\mathbf{Cx}_{\mathcal{S}}$ and $\mathbf{Sb}_{\mathcal{S}}(\Gamma, \Gamma)$, this is for every $\Gamma \in \mathcal{V}_\omega$ a function $\Gamma \rightarrow \Gamma$, which we can simply take to be the identity function:

$$\begin{aligned} \mathbf{id}_{\mathcal{S}} &: \prod_{\Gamma \in \mathcal{V}_\omega} \Gamma \rightarrow \Gamma \\ \mathbf{id}_{\mathcal{S}} \Gamma x &:= x \end{aligned}$$

Next, given any $\Gamma_0, \Gamma_1, \Gamma_2 \in \text{Cx}_{\mathcal{S}}$, $\gamma_1 \in \text{Sb}_{\mathcal{S}}(\Gamma_2, \Gamma_1)$, and $\gamma_0 \in \text{Sb}_{\mathcal{S}}(\Gamma_1, \Gamma_0)$ we must define the composite \mathcal{S} -substitution $\gamma_0 \circ_{\mathcal{S}} \gamma_1 \in \text{Sb}_{\mathcal{S}}(\Gamma_2, \Gamma_0)$, namely by function composition:

$$\begin{aligned} - \circ_{\mathcal{S}} - &: \prod_{\Gamma_0, \Gamma_1, \Gamma_2 \in \text{Cx}_{\mathcal{S}}} \text{Sb}_{\mathcal{S}}(\Gamma_1, \Gamma_0) \rightarrow \text{Sb}_{\mathcal{S}}(\Gamma_2, \Gamma_1) \rightarrow \text{Sb}_{\mathcal{S}}(\Gamma_2, \Gamma_0) \\ (\gamma_0 \circ_{\mathcal{S}} \gamma_1)(x) &:= \gamma_0(\gamma_1(x)) \end{aligned}$$

Notation 3.5.8. Starting with the above definition, we suppress unambiguous arguments for clarity: in this case, the \mathcal{S} -contexts $\Gamma_0, \Gamma_1, \Gamma_2$.

In the substitution calculus, identity and composition satisfy various equations, namely that composition is associative with identity as a left and right unit. We must therefore verify that our definitions of \mathcal{S} -identity and \mathcal{S} -composition validate the same equations:

Exercise 3.10. Verify the following equations:

- For all $\gamma \in \text{Sb}_{\mathcal{S}}(\Delta, \Gamma)$, $\text{id}_{\mathcal{S}} \circ_{\mathcal{S}} \gamma = \gamma = \gamma \circ_{\mathcal{S}} \text{id}_{\mathcal{S}}$.
- For all $\gamma_2 \in \text{Sb}_{\mathcal{S}}(\Gamma_3, \Gamma_2)$, $\gamma_1 \in \text{Sb}_{\mathcal{S}}(\Gamma_2, \Gamma_1)$, and $\gamma_0 \in \text{Sb}_{\mathcal{S}}(\Gamma_1, \Gamma_0)$, $\gamma_0 \circ_{\mathcal{S}} (\gamma_1 \circ_{\mathcal{S}} \gamma_2) = (\gamma_0 \circ_{\mathcal{S}} \gamma_1) \circ_{\mathcal{S}} \gamma_2$.

The empty context Next we define the empty \mathcal{S} -context $\mathbf{1}_{\mathcal{S}} \in \text{Cx}_{\mathcal{S}}$ and the terminal \mathcal{S} -substitution $!_{\mathcal{S}} \in \text{Sb}_{\mathcal{S}}(\Gamma, \mathbf{1}_{\mathcal{S}})$ for every $\Gamma \in \text{Cx}_{\mathcal{S}}$. Notably, although we call $\mathbf{1}$ the *empty* context, it is in fact interpreted as a *one-element* set.

$$\begin{aligned} \mathbf{1}_{\mathcal{S}} &\in \mathcal{V}_{\omega} \\ \mathbf{1}_{\mathcal{S}} &:= \{\star\} \end{aligned}$$

Remark 3.5.9. We write $\{\star\}$ to emphasize that it does not matter which one-element set in \mathcal{V}_{ω} we choose. The most natural concrete choice of one-element set is perhaps $\{\emptyset\}$, which we note is an element of \mathcal{V}_{ω} by axioms (1), (2) and (3) of Definition 3.5.1. \diamond

Exercise 3.11. In light of the definition of $\mathbf{1}_{\mathcal{S}}$ above, show that closed \mathcal{S} -types are just sets and closed \mathcal{S} -terms are just elements of those sets. To be precise, construct isomorphisms $\iota : \text{Ty}_{\mathcal{S}}(\mathbf{1}_{\mathcal{S}}) \cong \mathcal{V}_{\omega}$ and $\kappa_A : \text{Tm}_{\mathcal{S}}(\mathbf{1}_{\mathcal{S}}, A) \cong \iota(A)$ for all $A \in \text{Ty}_{\mathcal{S}}(\mathbf{1}_{\mathcal{S}})$.

The terminal \mathcal{S} -substitution into $\mathbf{1}_{\mathcal{S}}$ is the constant function returning \star .

$$\begin{aligned} !_{\mathcal{S}} &: \prod_{\Gamma \in \text{Cx}_{\mathcal{S}}} \text{Sb}_{\mathcal{S}}(\Gamma, \mathbf{1}_{\mathcal{S}}) \\ !_{\mathcal{S}}(x) &:= \star \end{aligned}$$

We have one equation to check before moving on.

Lemma 3.5.10. For all $\delta \in \text{Sb}_{\mathcal{S}}(\Gamma, \mathbf{1}_{\mathcal{S}})$, $\delta = !_{\mathcal{S}}$.

Proof. Unfolding definitions, we see that δ and $!_{\mathcal{S}}$ are both functions $\Gamma \rightarrow \{\star\}$. There is only one such function, so they must be equal. \square

Applying substitutions Applying an \mathcal{S} -substitution $\Delta \rightarrow \Gamma$ to an \mathcal{S} -type (resp., \mathcal{S} -term) in context Γ must produce an \mathcal{S} -type (resp., \mathcal{S} -term) in context Δ :

$$\begin{aligned} _[-]_{\mathcal{S}} &: \prod_{\Delta, \Gamma \in \mathbf{Cx}_{\mathcal{S}}} \prod_{\gamma \in \mathbf{Sb}_{\mathcal{S}}(\Delta, \Gamma)} \mathbf{T}_{\mathcal{S}}(\Gamma) \rightarrow \mathbf{T}_{\mathcal{S}}(\Delta) \\ _[-]_{\mathcal{S}} &: \prod_{\Delta, \Gamma \in \mathbf{Cx}_{\mathcal{S}}} \prod_{\gamma \in \mathbf{Sb}_{\mathcal{S}}(\Delta, \Gamma)} \prod_{A \in \mathbf{T}_{\mathcal{S}}(\Gamma)} \mathbf{Tm}_{\mathcal{S}}(\Gamma, A) \rightarrow \mathbf{Tm}_{\mathcal{S}}(\Delta, A[\gamma]_{\mathcal{S}}) \end{aligned}$$

Thankfully, the types of these operations are significantly more intimidating than their definitions. Unfolding definitions in the first line, we must take a function $\gamma : \Delta \rightarrow \Gamma$ and a function $A : \Gamma \rightarrow \mathcal{V}_{\omega}$ and produce a function $\Delta \rightarrow \mathcal{V}_{\omega}$, which is easily accomplished by composing A and γ . Substitution on terms is identical:

$$\begin{aligned} A[\gamma]_{\mathcal{S}} &:= A \circ \gamma \\ a[\gamma]_{\mathcal{S}} &:= a \circ \gamma \end{aligned}$$

The substitution calculus includes a number of equations governing $_[-]_{\mathcal{S}}$, namely that substituting by \mathbf{id} is the identity and substituting by a composite substitution is the same as a composition of substitutions; checking these for \mathcal{S} is again straightforward.

Exercise 3.12. Verify the following, where $\Gamma \in \mathbf{Cx}_{\mathcal{S}}$, $A \in \mathbf{T}_{\mathcal{S}}(\Gamma)$, and $a \in \mathbf{Tm}_{\mathcal{S}}(\Gamma, A)$:

- $A[\mathbf{id}_{\mathcal{S}}]_{\mathcal{S}} = A$.
- $a[\mathbf{id}_{\mathcal{S}}]_{\mathcal{S}} = a$.
- If $\gamma_1 \in \mathbf{Sb}_{\mathcal{S}}(\Gamma_2, \Gamma_1)$ and $\gamma_0 \in \mathbf{Sb}_{\mathcal{S}}(\Gamma_1, \Gamma)$, then $A[\gamma_0 \circ_{\mathcal{S}} \gamma_1]_{\mathcal{S}} = A[\gamma_0]_{\mathcal{S}}[\gamma_1]_{\mathcal{S}}$.
- If $\gamma_1 \in \mathbf{Sb}_{\mathcal{S}}(\Gamma_2, \Gamma_1)$ and $\gamma_0 \in \mathbf{Sb}_{\mathcal{S}}(\Gamma_1, \Gamma)$, then $a[\gamma_0 \circ_{\mathcal{S}} \gamma_1]_{\mathcal{S}} = a[\gamma_0]_{\mathcal{S}}[\gamma_1]_{\mathcal{S}}$.

Extending contexts The remaining operations of the substitution calculus are context extension $\Gamma.A$, substitution extension $\gamma.a$, the weakening substitution \mathbf{p} , and the variable term \mathbf{q} . We must start by defining the \mathcal{S} -interpretation of context extension, because it occurs in the types of all the other operations.

Recall from Sections 2.3 and 2.4.2 that substitutions into $\Gamma.A$ are roughly “pairs of a substitution into Γ and a term of type A .” More precisely, there is a natural isomorphism between substitutions $\gamma \in \mathbf{Sb}(\Delta, \Gamma.A)$ and pairs $(\gamma_0 \in \mathbf{Sb}(\Delta, \Gamma), a \in \mathbf{Tm}(\Delta, A[\gamma_0]))$. Unfolding \mathcal{S} -interpretations and setting $\Delta = \mathbf{1}_{\mathcal{S}}$, in light of Exercise 3.11 we see that elements of the set $\Gamma.\mathcal{S}A$ must be in bijection with pairs $(x_0 \in \Gamma, a \in A(x_0))$, and so we might as well take this as the definition of $\Gamma.\mathcal{S}A$.

$$\begin{aligned} _[\cdot]_{\mathcal{S}} &: \prod_{\Gamma \in \mathbf{Cx}_{\mathcal{S}}} \mathbf{T}_{\mathcal{S}}(\Gamma) \rightarrow \mathbf{Cx}_{\mathcal{S}} \\ \Gamma.\mathcal{S}A &:= \sum_{x \in \Gamma} A(x) \end{aligned}$$

We must be careful to check that this set is actually an element of $Cx_S = \mathcal{V}_\omega$, which follows from the closure of Grothendieck universes under indexed coproducts (Lemma 3.5.2).

Once again, to define \mathbf{p}_S , \mathbf{q}_S , and $_{.}S_{.}$ we must unfold their types, which will turn out to be significantly more intimidating than their definitions. Weakening, for example, is simply the first projection from Σ :

$$\begin{aligned} \mathbf{p}_S &: \prod_{\Gamma \in \mathcal{V}_\omega} \prod_{A \in \Gamma \rightarrow \mathcal{V}_\omega} (\sum_{x \in \Gamma} A(x)) \rightarrow \Gamma \\ \mathbf{p}_S(x, a) &:= x \end{aligned}$$

Similarly, variables and substitution extension are respectively the second projection and pairing operations of Σ . For any $\Delta, \Gamma \in \mathcal{V}_\omega$ and $A \in \Gamma \rightarrow \mathcal{V}_\omega$:

$$\begin{aligned} \mathbf{q}_S &: \prod_{p \in (\sum_{x \in \Gamma} A(x))} A(\mathbf{p}_S(p)) \\ \mathbf{q}_S(x, a) &:= a \\ \\ \text{_{.}S_{.}} &: \prod_{\gamma \in \Delta \rightarrow \Gamma} (\prod_{y \in \Delta} A(\gamma(y))) \rightarrow \Delta \rightarrow \sum_{x \in \Gamma} A(x) \\ (\gamma.Sa)(y) &:= (\gamma(y), a(y)) \end{aligned}$$

Exercise 3.13. Check that the types given above for \mathbf{p}_S , \mathbf{q}_S , and $_{.}S_{.}$ match the types given in Section 2.3, by unfolding the \mathcal{S} -interpretations given throughout this section.

The \mathcal{S} -interpretations of context extension, substitution extension, weakening, and variables as Σ , pairing, first projection, and second projection may in fact clarify the meaning of these operations in the substitution calculus. At any rate, it is straightforward to verify the necessary equations, which correspond to the β - and η -laws of Σ .

Lemma 3.5.11. *If $\Delta, \Gamma \in Cx_S$ and $A \in Ty_S(\Gamma)$, then:*

- *If $\gamma \in Sb_S(\Delta, \Gamma)$ and $a \in Tm_S(\Delta, A[\gamma]_S)$, then $\mathbf{p}_S \circ_S (\gamma.Sa) = \gamma$.*
- *If $\gamma \in Sb_S(\Delta, \Gamma)$ and $a \in Tm_S(\Delta, A[\gamma]_S)$, then $\mathbf{q}_S[\gamma.Sa] = a$.*
- *If $\gamma \in Sb_S(\Delta, \Gamma.SA)$ then $\gamma = (\mathbf{p}_S \circ_S \gamma).S(\mathbf{q}_S[\gamma]_S)$.*

Proof. These all follow essentially by definition. For the first equation, fix $\gamma : \Delta \rightarrow \Gamma$ and $a \in \prod_{y \in \Delta} A(\gamma(y))$; we must show $\pi_1 \circ (\lambda y \rightarrow (\gamma(y), a(y))) = \gamma$. Because both sides are functions, it suffices to check that they agree on all $y \in \Delta$, and indeed both produce $\gamma(y)$ when applied to y . For the second equation we must show $\pi_2 \circ (\lambda y \rightarrow (\gamma(y), a(y))) = a$, which again follows by applying both sides to $y \in \Delta$.

For the third equation, fix $\gamma : \Delta \rightarrow \sum_{x \in \Gamma} A(x)$ and show $\gamma = \lambda y \rightarrow (\pi_1(\gamma(y)), \pi_2(\gamma(y)))$. This follows by applying both sides to $y \in \Delta$ and noting that $\gamma(y) \in \sum_{x \in \Gamma} A(x)$ is by definition of the form (x_0, a) . \square

The reader should now verify that we have provided an \mathcal{S} -interpretation of every rule of the substitution calculus, covering the first section of Appendix A.

Notation 3.5.12. We note that we can safely reuse notations from Chapter 2 for their \mathcal{S} counterparts. In particular, following Exercise 2.4, we write $\gamma.\mathcal{S}A$ for $(\gamma \circ_{\mathcal{S}} \mathbf{p}_{\mathcal{S}}).\mathcal{S}(\mathbf{q}_{\mathcal{S}})_{\mathcal{S}}$.

3.5.3 The type-theoretic connectives of sets

Now that we have defined the \mathcal{S} -interpretation of the basic structure of type theory, we can extend \mathcal{S} with any connectives of our choice. Unlike the operations considered in Section 3.5.2, the connectives of type theory are (generally) defined independently of one another, allowing us to model them in a modular fashion. We consider some representative cases, namely, the \mathcal{S} -interpretations of Π -types, **Eq**-types, **Void**, **Bool**, and U_0 .

Π -types Taking advantage of the compact representation of the rules of Π -types introduced in Section 2.4.2, the \mathcal{S} -interpretation of Π -types consists of an \mathcal{S} -type-forming operation and a family of isomorphisms of sets:

$$\begin{aligned} \Pi_{\mathcal{S}} &: \prod_{\Gamma \in \mathbf{C}x_{\mathcal{S}}} (\sum_{A \in \mathbf{T}y_{\mathcal{S}}(\Gamma)} \mathbf{T}y_{\mathcal{S}}(\Gamma.\mathcal{S}A)) \rightarrow \mathbf{T}y_{\mathcal{S}}(\Gamma) \\ \iota_{\mathcal{S}} &: \prod_{\Gamma \in \mathbf{C}x_{\mathcal{S}}} \prod_{A \in \mathbf{T}y_{\mathcal{S}}(\Gamma)} \prod_{B \in \mathbf{T}y_{\mathcal{S}}(\Gamma.\mathcal{S}A)} \mathbf{T}m_{\mathcal{S}}(\Gamma, \Pi_{\mathcal{S}} \Gamma (A, B)) \cong \mathbf{T}m_{\mathcal{S}}(\Gamma.\mathcal{S}A, B) \end{aligned}$$

subject to the following equations expressing their naturality in $\Gamma \in \mathbf{C}x_{\mathcal{S}}$:

$$\begin{aligned} (\Pi_{\mathcal{S}} \Gamma (A, B))[\gamma]_{\mathcal{S}} &= \Pi_{\mathcal{S}} \Delta (A[\gamma]_{\mathcal{S}}, B[\gamma.\mathcal{S}A]_{\mathcal{S}}) & (\gamma \in \mathbf{S}b_{\mathcal{S}}(\Delta, \Gamma)) \\ (\iota_{\mathcal{S}} \Gamma A B f)[\gamma.\mathcal{S}A]_{\mathcal{S}} &= \iota_{\mathcal{S}} \Delta (A[\gamma]_{\mathcal{S}}) (B[\gamma.\mathcal{S}A]_{\mathcal{S}}) (f[\gamma]_{\mathcal{S}}) & (\gamma \in \mathbf{S}b_{\mathcal{S}}(\Delta, \Gamma)) \end{aligned}$$

To get a handle on the situation, let us consider the types of $\Pi_{\mathcal{S}}$ and $\iota_{\mathcal{S}}$ when specialized to the empty context $1_{\mathcal{S}}$, and simplified along the isomorphisms of Exercise 3.11:

$$\begin{aligned} \Pi_{\mathcal{S}} 1_{\mathcal{S}} &: (\sum_{A \in \mathcal{V}_{\omega}} (A \rightarrow \mathcal{V}_{\omega})) \rightarrow \mathcal{V}_{\omega} \\ \iota_{\mathcal{S}} 1_{\mathcal{S}} &: \prod_{A \in \mathcal{V}_{\omega}} \prod_{B \in A \rightarrow \mathcal{V}_{\omega}} \Pi_{\mathcal{S}} 1_{\mathcal{S}} (A, B) \cong \prod_{a \in A} B(a) \end{aligned}$$

That is, in the empty context, for any $A \in \mathcal{V}_{\omega}$ and $B : A \rightarrow \mathcal{V}_{\omega}$ we must choose a set $\Pi_{\mathcal{S}} 1_{\mathcal{S}} (A, B) \in \mathcal{V}_{\omega}$ to serve as the \mathcal{S} - Π -type of A and B , and this set must be isomorphic to the set-theoretic indexed product $\prod_{a \in A} B(a)$.

The situation for arbitrary contexts is essentially the same, except that all three of A , B , and their \mathcal{S} - Π -type are additionally indexed by a set Γ . We define $\Pi_{\mathcal{S}}$ as follows:

$$\Pi_{\mathcal{S}} \Gamma (A, B) x := \prod_{a \in A(x)} B(x, a) \quad (x \in \Gamma)$$

noting that $B : (\sum_{x \in \Gamma} A(x)) \rightarrow \mathcal{V}_{\omega}$ by the definition of $\Gamma.\mathcal{S}A$ in Section 3.5.2. Finally, we must verify that our definition $\prod_{a \in A(x)} B(x, a) \in \mathcal{V}_{\omega}$, which indeed holds by Lemma 3.5.2.

Lemma 3.5.13. $\Pi_{\mathcal{S}}$ is natural in Γ , i.e., $(\Pi_{\mathcal{S}} \Gamma (A, B))[\gamma]_{\mathcal{S}} = \Pi_{\mathcal{S}} \Delta (A[\gamma]_{\mathcal{S}}, B[\gamma.\mathcal{S}A]_{\mathcal{S}})$ in $\text{Ty}_{\mathcal{S}}(\Delta)$ for any $\gamma \in \text{Sb}_{\mathcal{S}}(\Delta, \Gamma)$.

Proof. Unfolding the operations of the substitution calculus, we must show:

$$(\Pi_{\mathcal{S}} \Gamma (A, B)) \circ \gamma = \Pi_{\mathcal{S}} \Delta (A \circ \gamma, \lambda(y, a) \rightarrow B(\gamma(y), a))$$

These are both functions $\Delta \rightarrow \mathcal{V}_{\omega}$, so it suffices to check that they agree on all $y \in \Delta$:

$$\begin{aligned} & ((\Pi_{\mathcal{S}} \Gamma (A, B)) \circ \gamma)(y) \\ &= \Pi_{\mathcal{S}} \Gamma (A, B) (\gamma(y)) \\ &= \prod_{a \in A(\gamma(y))} B(\gamma(y), a) \\ &= \prod_{a \in (A \circ \gamma)(y)} (\lambda(y, a) \rightarrow B(\gamma(y), a))(y, a) \\ &= \Pi_{\mathcal{S}} \Delta (A \circ \gamma, \lambda(y, a) \rightarrow B(\gamma(y), a)) y \quad \square \end{aligned}$$

As for the isomorphism $\iota_{\mathcal{S}}$, unfolding definitions we must construct:

$$\iota_{\mathcal{S}} : \prod_{\Gamma \in \mathcal{V}_{\omega}} \prod_{A \in \Gamma \rightarrow \mathcal{V}_{\omega}} \prod_{B \in (\sum_{x \in \Gamma} A(x)) \rightarrow \mathcal{V}_{\omega}} (\prod_{x \in \Gamma} \prod_{a \in A(x)} B(x, a)) \cong \prod_{p \in (\sum_{x \in \Gamma} A(x))} B(p)$$

Fixing Γ, A, B , this isomorphism is simply the dependent (un)currying isomorphism $(x : \Gamma) \rightarrow (a : A(x)) \rightarrow B(x, a) \cong (p : \sum_{x \in \Gamma} A(x)) \rightarrow B(p)$, defined as follows:

$$\begin{aligned} \iota_{\mathcal{S}} \Gamma A B f (x, a) &:= f x a \\ \iota_{\mathcal{S}}^{-1} \Gamma A B g x a &:= g (x, a) \end{aligned}$$

Exercise 3.14. Verify that $\iota_{\mathcal{S}}$ and $\iota_{\mathcal{S}}^{-1}$ are inverses.

Exercise 3.15. Verify that $\iota_{\mathcal{S}}$ is natural in Γ . (Hint: show that

$$(\iota_{\mathcal{S}} \Gamma A B f) \circ (\lambda(y, a) \rightarrow (\gamma(y), a)) = \iota_{\mathcal{S}} \Delta (A \circ \gamma) (\lambda(y, a) \rightarrow B(\gamma(y), a)) (f \circ \gamma)$$

for any $f \in \prod_{x \in \Gamma} \prod_{a \in A(x)} B(x, a)$ and $\gamma \in \text{Sb}_{\mathcal{S}}(\Delta, \Gamma)$, by showing that they agree on all $(y, a) \in \sum_{y \in \Delta} A(\gamma(y))$.)

Eq-types The \mathcal{S} -interpretation of extensional equality types is analogous to that of Π -types. Following Section 2.4.4, we must define an \mathcal{S} -type-forming operation and a family of isomorphisms of sets, both natural in Γ :

$$\begin{aligned} \text{Eq}_{\mathcal{S}} &: \prod_{\Gamma \in \text{Cx}_{\mathcal{S}}} (\sum_{A \in \text{Ty}_{\mathcal{S}}(\Gamma)} \text{Tm}_{\mathcal{S}}(\Gamma, A) \times \text{Tm}_{\mathcal{S}}(\Gamma, A)) \rightarrow \text{Ty}_{\mathcal{S}}(\Gamma) \\ \iota_{\mathcal{S}} &: \prod_{\Gamma \in \text{Cx}_{\mathcal{S}}} \prod_{A \in \text{Ty}_{\mathcal{S}}(\Gamma)} \prod_{a, b \in \text{Tm}_{\mathcal{S}}(\Gamma, A)} \text{Tm}_{\mathcal{S}}(\Gamma, \text{Eq}_{\mathcal{S}} \Gamma (A, a, b)) \cong \{\star \mid a = b\} \end{aligned}$$

We define $\mathbf{Eq}_{\mathcal{S}} \Gamma (A, a, b)$ to be the Γ -indexed family of sets that maps $x \in \Gamma$ to a one-element set when $a(x) = b(x) \in A(x)$, and an empty set otherwise.

$$\mathbf{Eq}_{\mathcal{S}} \Gamma (A, a, b) x := \{\star \mid a(x) = b(x)\}$$

To define $\iota_{\mathcal{S}}$, we note that \mathcal{S} -terms $e \in \mathbf{Tm}_{\mathcal{S}}(\Gamma, \mathbf{Eq}_{\mathcal{S}} \Gamma (A, a, b))$ are constant functions sending every $x \in \Gamma$ to the unique element \star . In particular, the existence of such an e implies that $a(x) = b(x)$ for all $x \in \Gamma$, and any two such terms e, e' must agree on all $x \in \Gamma$ and thus be equal. Thus:

$$\begin{aligned} \iota_{\mathcal{S}} \Gamma A a b e &:= \star \\ \iota_{\mathcal{S}}^{-1} \Gamma A a b \star x &:= \star \end{aligned}$$

Exercise 3.16. Verify that $\iota_{\mathcal{S}}$ and $\iota_{\mathcal{S}}^{-1}$ are inverses.

Exercise 3.17. State and prove the naturality equations for $\mathbf{Eq}_{\mathcal{S}}$ and $\iota_{\mathcal{S}}$. (Hint: reference the naturality equations in Section 2.4.4, and unfold definitions.)

The empty type Our next type \mathbf{Void} is defined not by a mapping-in property but a mapping-out property. However, as discussed in Section 2.5.1, it can nevertheless be axiomatized as a natural type-forming operation with a natural family of isomorphisms:

$$\begin{aligned} \mathbf{Void}_{\mathcal{S}} &: \prod_{\Gamma \in \mathcal{C}X_{\mathcal{S}}} \mathbf{Ty}_{\mathcal{S}}(\Gamma) \\ \rho_{\mathcal{S}} &: \prod_{\Gamma \in \mathcal{C}X_{\mathcal{S}}} \prod_{A \in \mathbf{Ty}_{\mathcal{S}}(\Gamma, \mathcal{S}\mathbf{Void}_{\mathcal{S}})} \mathbf{Tm}_{\mathcal{S}}(\Gamma, \mathcal{S}\mathbf{Void}_{\mathcal{S}}, A) \cong \{\star\} \end{aligned}$$

Given that \mathbf{Void} is called the empty type, it is perhaps unsurprising that \mathcal{S} interprets it as the empty set, regarded as a constant family over $\Gamma \in \mathcal{V}_{\omega}$ and $x \in \Gamma$.

$$\mathbf{Void}_{\mathcal{S}} \Gamma x := \emptyset$$

Elements of the \mathcal{S} -context $\Gamma, \mathcal{S}\mathbf{Void}_{\mathcal{S}}$ are pairs of $x \in \Gamma$ and $y \in \mathbf{Void}_{\mathcal{S}} x$, but the latter set is defined to be empty, so no such pairs exist and $\Gamma, \mathcal{S}\mathbf{Void}_{\mathcal{S}} = \emptyset$. Accordingly, \mathcal{S} -terms $f \in \mathbf{Tm}_{\mathcal{S}}(\Gamma, \mathcal{S}\mathbf{Void}_{\mathcal{S}}, A)$ are (dependent) functions out of an empty set. As discussed in Section 2.5.1, there is exactly one such function for every A , and this is precisely the content of the isomorphism $\rho_{\mathcal{S}}$.

$$\rho_{\mathcal{S}} \Gamma A a := \star$$

Exercise 3.18. Complete the \mathcal{S} -interpretation of \mathbf{Void} : verify that $\rho_{\mathcal{S}}$ is an isomorphism, and prove the naturality equations for $\mathbf{Void}_{\mathcal{S}}$ and $\rho_{\mathcal{S}}$, following Section 2.5.1.

Booleans Like **Void**, the booleans are also defined by a mapping-out property. Recalling Section 2.5.2, the specification of **Bool** has three components, the first two being a natural type-former and two natural term-formers:

$$\begin{aligned} \mathbf{Bool}_S &: \prod_{\Gamma \in \mathcal{C}_{X_S}} \mathbf{Ty}_S(\Gamma) \\ \mathbf{true}_S, \mathbf{false}_S &: \prod_{\Gamma \in \mathcal{C}_{X_S}} \mathbf{Tm}_S(\Gamma, \mathbf{Bool}_S) \end{aligned}$$

The third component is once again a natural isomorphism, but unlike the previous examples in which the two directions of the isomorphism encode introduction and elimination, here the forward map is fixed by the choice of \mathbf{true}_S and \mathbf{false}_S , and the reverse map expresses the principle that maps out of **Bool** are determined by their instantiations at \mathbf{true} and \mathbf{false} . Writing $\rho \Gamma A$ for the map which sends $a \in \mathbf{Tm}_S(\Gamma, \mathcal{S}\mathbf{Bool}_S, A)$ to the pair of \mathcal{S} -terms $(a[\mathbf{id}_S.\mathcal{S}\mathbf{true}_S]_S, a[\mathbf{id}_S.\mathcal{S}\mathbf{false}_S]_S)$, we require ρ to be an isomorphism.

$$\begin{aligned} \rho &: \prod_{\Gamma \in \mathcal{C}_{X_S}} \prod_{A \in \mathbf{Ty}_S(\Gamma, \mathcal{S}\mathbf{Bool}_S)} \mathbf{Tm}_S(\Gamma, \mathcal{S}\mathbf{Bool}_S, A) \cong \\ &\quad \mathbf{Tm}_S(\Gamma, A[\mathbf{id}_S.\mathcal{S}\mathbf{true}_S]_S) \times \mathbf{Tm}_S(\Gamma, A[\mathbf{id}_S.\mathcal{S}\mathbf{false}_S]_S) \\ \rho \Gamma A a &:= (a[\mathbf{id}_S.\mathcal{S}\mathbf{true}_S]_S, a[\mathbf{id}_S.\mathcal{S}\mathbf{false}_S]_S) \end{aligned}$$

We can define \mathbf{Bool}_S to be any fixed two-element set $\{\mathbf{true}_S, \mathbf{false}_S\}$, regarded as a constant family over $\Gamma \in \mathcal{V}_\omega$ and $x \in \Gamma$.

$$\begin{aligned} \mathbf{Bool}_S \Gamma x &:= \{0, 1\} \\ \mathbf{true}_S \Gamma x &:= 1 \\ \mathbf{false}_S \Gamma x &:= 0 \end{aligned}$$

Exercise 3.19. State and prove the naturality equations for \mathbf{Bool}_S , \mathbf{true}_S , and \mathbf{false}_S .

It remains only to check that $\rho \Gamma A$ is indeed an isomorphism.

Lemma 3.5.14. *The map $a \mapsto (a[\mathbf{id}_S.\mathcal{S}\mathbf{true}_S]_S, a[\mathbf{id}_S.\mathcal{S}\mathbf{false}_S]_S)$ is an isomorphism*

$$\mathbf{Tm}_S(\Gamma, \mathcal{S}\mathbf{Bool}_S, A) \cong \mathbf{Tm}_S(\Gamma, A[\mathbf{id}_S.\mathcal{S}\mathbf{true}_S]_S) \times \mathbf{Tm}_S(\Gamma, A[\mathbf{id}_S.\mathcal{S}\mathbf{false}_S]_S)$$

Proof. Unfolding definitions, the \mathcal{S} -context $\Gamma.\mathcal{S}\mathbf{Bool}_S$ is the set $\Gamma \times \{0, 1\}$, so \mathcal{S} -types $A \in \mathbf{Ty}_S(\Gamma, \mathcal{S}\mathbf{Bool}_S)$ are families of sets $\Gamma \times \{0, 1\} \rightarrow \mathcal{V}_\omega$, and \mathcal{S} -terms $a \in \mathbf{Tm}_S(\Gamma, \mathcal{S}\mathbf{Bool}_S, A)$ are dependent functions $\prod_{p \in \Gamma \times \{0, 1\}} A(p)$. But

$$\begin{aligned} &\prod_{p \in \Gamma \times \{0, 1\}} A(p) \\ \cong &\prod_{x \in \Gamma} \prod_{b \in \{0, 1\}} A(x, b) \\ \cong &\prod_{x \in \Gamma} A(x, 1) \times A(x, 0) \\ \cong &(\prod_{x \in \Gamma} A(x, 1)) \times (\prod_{x \in \Gamma} A(x, 0)) \end{aligned}$$

where the forward composite map is $a \mapsto ((\lambda x \rightarrow a(x, 1)), (\lambda x \rightarrow a(x, 0)))$. Unfolding definitions, this is precisely the map we wanted to show is an isomorphism. \square

Universes The final connective we discuss is \mathbf{U} , a “type of types” whose terms $\Gamma \vdash a : \mathbf{U}$ decode to types $\Gamma \vdash \mathbf{El}(a)$ type. As we saw in Section 2.6, universe types require far more rules than the other connectives: type theory has a countably infinite hierarchy of universes $\mathbf{U} = \mathbf{U}_0 : \mathbf{U}_1 : \mathbf{U}_2 : \dots$, each closed under codes for every type-former and satisfying definitional equalities involving \mathbf{El} , with **lift** operations between these universes commuting with all the aforementioned operations. In addition, the \mathcal{S} -interpretation of \mathbf{U} as a “set of sets” will force us to confront some set-theoretic technicalities.

The good news is that all of this structure will fall quite neatly into place. The astute reader may have noticed that Axiom 3.5.6 postulates an infinite hierarchy of Grothendieck universes $\mathcal{V}_0 \in \dots \in \mathcal{V}_\omega$ of which we have only used \mathcal{V}_ω thus far; the remaining \mathcal{V}_i serve as the \mathcal{S} -interpretations of the type-theoretic universe hierarchy.

Let us begin by defining $(\mathbf{U}_0)_\mathcal{S} = \mathbf{U}_\mathcal{S}$ and $(\mathbf{El}_0)_\mathcal{S} = \mathbf{El}_\mathcal{S}$:

$$\begin{aligned} \mathbf{U}_\mathcal{S} &: \prod_{\Gamma \in \mathcal{V}_\omega} \mathbf{Ty}_\mathcal{S}(\Gamma) \\ \mathbf{U}_\mathcal{S} \Gamma x &:= \mathcal{V}_0 \\ \mathbf{El}_\mathcal{S} &: \prod_{\Gamma \in \mathcal{V}_\omega} \mathbf{Tm}_\mathcal{S}(\Gamma, \mathbf{U}_\mathcal{S}) \rightarrow \mathbf{Ty}_\mathcal{S}(\Gamma) \\ \mathbf{El}_\mathcal{S} \Gamma c &:= c \end{aligned}$$

To make sense of the last definition, we note that $\mathbf{El}_\mathcal{S} \Gamma : (\Gamma \rightarrow \mathcal{V}_0) \rightarrow (\Gamma \rightarrow \mathcal{V}_\omega)$. By our hypothesis $\mathcal{V}_0 \in \mathcal{V}_\omega$ and Lemma 3.5.2, $\mathcal{V}_0 \subseteq \mathcal{V}_\omega$, so in particular $(\Gamma \rightarrow \mathcal{V}_0) \subseteq (\Gamma \rightarrow \mathcal{V}_\omega)$.

Exercise 3.20. State and prove the naturality equations for $\mathbf{U}_\mathcal{S}$ and $\mathbf{El}_\mathcal{S}$.

Following Section 2.6.2, the \mathcal{S} -interpretation of \mathbf{U} must include codes for $\mathbf{\Pi}$ -types:

$$\mathbf{pi}_\mathcal{S} : \prod_{\Gamma \in \mathcal{C}_{X_\mathcal{S}}} (\sum_{A \in \mathbf{Tm}_\mathcal{S}(\Gamma, \mathbf{U}_\mathcal{S})} \mathbf{Tm}_\mathcal{S}(\Gamma, \mathcal{S}\mathbf{El}_\mathcal{S}(A), \mathbf{U}_\mathcal{S})) \rightarrow \mathbf{Tm}_\mathcal{S}(\Gamma, \mathbf{U}_\mathcal{S})$$

satisfying a naturality equation as well as the following equation in $\Gamma \rightarrow \mathcal{V}_\omega$:

$$\mathbf{El}_\mathcal{S} \Gamma (\mathbf{pi}_\mathcal{S} \Gamma (A, B)) = \mathbf{\Pi}_\mathcal{S} \Gamma (\mathbf{El}_\mathcal{S} \Gamma A, \mathbf{El}_\mathcal{S} \Gamma B)$$

Because $\mathbf{El}_\mathcal{S} \Gamma$ is just the inclusion $(\Gamma \rightarrow \mathcal{V}_0) \subseteq (\Gamma \rightarrow \mathcal{V}_\omega)$, we can simply take the above equation as a *definition*—setting $\mathbf{pi}_\mathcal{S} \Gamma (A, B) := \mathbf{\Pi}_\mathcal{S} \Gamma (A, B)$ —as long as we prove that the right-hand side lands inside of $\Gamma \rightarrow \mathcal{V}_0$ when A and B are pointwise \mathcal{V}_0 -small.

Lemma 3.5.15. *If $\Gamma \in \mathcal{V}_\omega$, $A \in \Gamma \rightarrow \mathcal{V}_0$, and $B \in (\sum_{x \in \Gamma} A(x)) \rightarrow \mathcal{V}_0$, then*

$$(\prod_{x \in \Gamma} \prod_{a \in A(x)} B(x, a)) \in \Gamma \rightarrow \mathcal{V}_0$$

Proof. Note that this statement refines a similar observation in our construction of \mathcal{S} - $\mathbf{\Pi}$ -types, in which all the \mathcal{V}_0 are replaced by \mathcal{V}_ω . The proof is identical: because \mathcal{V}_0 is a Grothendieck universe, Lemma 3.5.2 implies that $\prod_{a \in A(x)} B(x, a) \in \mathcal{V}_0$ for all $x \in \Gamma$. \square

The naturality equation for \mathbf{pi}_S then follows immediately from the naturality of $\mathbf{\Pi}_S$. The codes for other connectives proceed identically, using the fact that \mathcal{V}_0 is closed under every relevant construction. For the remainder of the universe hierarchy, we define $(\mathbf{U}_i)_S \Gamma x := \mathcal{V}_i$ and check that $(\mathbf{El}_i)_S$ and $(\mathbf{lift}_i)_S$ are subset inclusions.

3.5.4 Using the set model

We finally arrive at the main result of this section.

Theorem 3.5.16. *\mathcal{S} is a model of extensional type theory.*

Although extensional type theory is often considered an *alternative* to set theory, the fact that \mathcal{S} allows us to reduce questions about type theory to questions about sets makes the set model one of the most powerful tools for studying the properties of type theory. In Section 3.6, we appeal to \mathcal{S} in two proofs that equality in extensional type theory is undecidable; in the remainder of this section, we will quickly rattle off several other corollaries of Theorem 3.5.16, starting with the consistency of type theory (Theorem 3.4.8).

Proof of Theorem 3.4.8. To show that type theory is consistent, by Theorem 3.4.7 it suffices to exhibit a model \mathcal{M} in which $\mathbf{Tm}_{\mathcal{M}}(\mathbf{1}_{\mathcal{M}}, \mathbf{Void}_{\mathcal{M}})$ is empty. Choosing $\mathcal{M} = \mathcal{S}$, by Exercise 3.11 we have $\mathbf{Tm}_{\mathcal{S}}(\mathbf{1}_{\mathcal{S}}, \mathbf{Void}_{\mathcal{S}}) \cong \mathbf{Void}_{\mathcal{S}} \mathbf{1}_{\mathcal{S}} \star := \emptyset$. \square

More generally, \mathcal{S} tells us that any term in extensional type theory—that is, in its syntactic model \mathcal{T} (Definition 3.4.4)—gives rise to a corresponding function of sets. On the one hand, this lets us construct functions on sets by writing down terms in type theory; on the other hand, we can *disprove* the existence of terms by showing that their image under the \mathcal{S} -interpretation does not exist, as we just did in the proof of consistency.

Lemma 3.5.17. *Within type theory, there are no injective functions $(\mathbf{Nat} \rightarrow \mathbf{Nat}) \rightarrow \mathbf{Nat}$; that is, there are no closed terms of type*

$$\sum_{f:(\mathbf{Nat} \rightarrow \mathbf{Nat}) \rightarrow \mathbf{Nat}} (g_1, g_2 : \mathbf{Nat} \rightarrow \mathbf{Nat}) \rightarrow f(g_1) = f(g_2) \rightarrow g_1 = g_2$$

Proof. Unfolding definitions, the image of such a term under \mathcal{S} is a pair whose first projection is an ordinary set-theoretic function $f : (\mathbb{N} \rightarrow \mathbb{N}) \rightarrow \mathbb{N}$, and whose second projection is a three-argument function that takes two functions $g_1, g_2 : \mathbb{N} \rightarrow \mathbb{N}$ and $x \in \{\star \mid f(g_1) = f(g_2)\}$, and returns $\{\star \mid g_1 = g_2\}$. In particular, although the second projection is unique when it exists, it exists only when f is injective. But $\mathbb{N} \rightarrow \mathbb{N}$ is uncountable, so there can be no injective functions from it to \mathbb{N} . \square

Remark 3.5.18. This argument does not go through if we restrict attention to the syntactic model, because the set $\mathbf{Tm}(\mathbf{1}, \mathbf{\Pi}(\mathbf{Nat}, \mathbf{Nat}))$ of closed terms of type $\mathbf{Nat} \rightarrow \mathbf{Nat}$ is *countable*: it is a quotient of a subset of finite derivation trees, which are countable. \diamond

Theorem 3.5.19. *Extensional type theory does not have injective Π -types (Definition 3.2.8).*

Proof. Using equality reflection and universes, the following judgment holds:

$$1.\text{Eq}(\mathbf{U}, \text{pi}(\text{unit}, \text{void}), \text{pi}(\text{bool}, \text{void})) \vdash \Pi(\mathbf{Unit}, \mathbf{Void}) = \Pi(\mathbf{Bool}, \mathbf{Void}) \text{ type}$$

If extensional type theory had injective Π -types, this would imply:

$$1.\text{Eq}(\mathbf{U}, \text{pi}(\text{unit}, \text{void}), \text{pi}(\text{bool}, \text{void})) \vdash \mathbf{Unit} = \mathbf{Bool} \text{ type}$$

This implies in particular that **true** and **false** are elements of **Unit** in this context. By the η rule for **Unit** this implies that **true** = **false** in this context, and hence by Theorem 2.6.3,

$$1.\text{Eq}(\mathbf{U}, \text{pi}(\text{unit}, \text{void}), \text{pi}(\text{bool}, \text{void})) \vdash \text{tt} : \mathbf{Void}$$

The \mathcal{S} -interpretation of the above context is a set with one element if $\Pi_{\mathcal{S}}(\mathbf{Unit}_{\mathcal{S}}, \mathbf{Void}_{\mathcal{S}})$ and $\Pi_{\mathcal{S}}(\mathbf{Bool}_{\mathcal{S}}, \mathbf{Void}_{\mathcal{S}})$ are equal sets, which is indeed the case because both are \emptyset . Thus the \mathcal{S} -interpretation of the above term must be a function from a one-element set to \emptyset , which does not exist. We conclude that there is no such term, and thus extensional type theory does not have injective Π -types. \square

Finally, recall that all of the constructions in this section have assumed an $(\omega + 1)$ -hierarchy of Grothendieck universes $\mathcal{V}_0 \in \dots \in \mathcal{V}_{\omega}$ (Axiom 3.5.6): we use \mathcal{V}_{ω} to model contexts and types, and smaller \mathcal{V}_i to model \mathbf{U}_i . In general, we need $n + 1$ Grothendieck universes to model a type theory with n universes.

Theorem 3.5.20. *An $(n + 1)$ -hierarchy of Grothendieck universes $\mathcal{V}_0 \in \dots \in \mathcal{V}_n$ suffices to construct a set-theoretic model of extensional type theory with n universes $\mathbf{U}_0 : \dots : \mathbf{U}_{n-1}$.*

3.6 Equality in extensional type theory is undecidable

In this section we present two proofs that term equality in extensional type theory is undecidable, and hence extensional type theory does not admit a normalization structure by Exercise 3.3. The first proof, due to Castellan, Clairambault, and Dybjer [CCD17], is conceptually straightforward but requires an appeal to the set-theoretic model (Section 3.5). The second proof, due to Hofmann [Hof95a], requires only the assumption that extensional type theory is consistent (Theorem 3.4.8), but is more complex, requiring the machinery of recursively inseparable sets. Both of these ideas arise with some frequency in the metatheory of type theory, so we cover both proofs in some detail.

3.6.1 The first proof: deciding equality of SK terms

The strategy of our first proof is to exhibit a context Γ_{SK} and an encoding $\llbracket - \rrbracket$ of terms of the SK combinator calculus into type-theoretic terms in context Γ_{SK} , such that two SK terms are convertible if and only if their encodings are judgmentally equal. Because convertibility of SK terms is undecidable, judgmental equality is as well.

Recall that the SK combinator calculus is an extremely minimal Turing-complete language generated by application and two combinators named S and K :

$$\text{Combinators } x := S \mid K \mid x x$$

Combinators compute according to the following rewriting system \mapsto . We say that two combinators are *convertible*, written $x \sim y$, if they are related by the reflexive, symmetric, and transitive closure of \mapsto .

$$\frac{}{S x y z \mapsto (x z) (y z)} \quad \frac{}{K x y \mapsto x} \quad \frac{x \mapsto x'}{x y \mapsto x' y} \quad \frac{y \mapsto y'}{x y \mapsto x y'}$$

We define the following context, written in Agda-style notation:

$$\begin{aligned} \Gamma_{SK} := & \mathbf{1}, \\ & A : \mathbf{U}, \\ & _ \bullet _ : A \rightarrow A \rightarrow A, \\ & s : A, \\ & k : A, \\ & e_1 : (a b : A) \rightarrow \mathbf{Eq}(A, (k \bullet a) \bullet b, a), \\ & e_2 : (a b c : A) \rightarrow \mathbf{Eq}(A, ((s \bullet a) \bullet b) \bullet c, (a \bullet c) \bullet (b \bullet c)) \end{aligned}$$

Writing Λ for the set of SK combinator terms, we can straightforwardly define a function $\llbracket - \rrbracket : \Lambda \rightarrow \text{Tm}(\Gamma_{SK}, A)$ by sending application, S , and K to \bullet , s , and k respectively, and this function respects convertibility of combinators.

Lemma 3.6.1. *There is a function $\llbracket - \rrbracket : \Lambda \rightarrow \text{Tm}(\Gamma_{SK}, A)$ such that $x \sim y \implies \llbracket x \rrbracket = \llbracket y \rrbracket$.*

Exercise 3.21. The context Γ_{SK} only includes two of the four generating rules of \mapsto . Why haven't we included the other two, or reflexivity, symmetry, or transitivity?

Lemma 3.6.1 implies that term equality is sound for an undecidable problem, but this does not yet imply that term equality is undecidable; it is possible, for example, that *all* terms in the image of $\llbracket - \rrbracket$ are equal. To complete our proof, we must observe that term equality is also *complete* for convertibility; we argue this by using the set-theoretic model of type theory to recover the convertibility class of x from the term $\llbracket x \rrbracket$.

Theorem 3.6.2. *If $\llbracket x \rrbracket = \llbracket y \rrbracket$ then $x \sim y$.*

Proof. Let us write $f : \mathcal{T} \rightarrow \mathcal{S}$ for the homomorphism from the syntactic model \mathcal{T} to the set-theoretic model \mathcal{S} . This homomorphism interprets syntactic contexts Γ as sets $C_{X_f}(\Gamma)$, syntactic types $A \in \text{Ty}(\Gamma)$ as $C_{X_f}(\Gamma)$ -indexed families of sets, and syntactic context extensions as indexed coproducts of those families. (See Section 3.5 for more details.)

Unwinding definitions, elements of $C_{X_f}(\Gamma_{SK})$ are “SK-algebras,” or dependent tuples of a set along with application, S , and K operations satisfying the convertibility axioms. Combinators modulo convertibility form such an algebra in the evident way; writing $[x]$ for the convertibility equivalence class of $x \in \Lambda$, we have

$$\gamma_{SK} := (\Lambda/\sim, (\lambda[x] [y] \rightarrow [x y]), [s], [k], \star, \star) \in C_{X_f}(\Gamma_{SK})$$

Homomorphisms of models respect equality, so from $\llbracket x \rrbracket = \llbracket y \rrbracket \in \text{Tm}(\Gamma_{SK}, A)$ we see that these terms are interpreted in \mathcal{S} as equal dependent functions $\prod_{(A, \dots) : C_{X_f}(\Gamma_{SK})} A$, and in particular, applying these functions to γ_{SK} produces two equal elements of Λ/\sim . We can prove by induction on combinators that for any $z \in \Lambda$ this procedure recovers z up to convertibility (*i.e.*, sends $\llbracket z \rrbracket$ to $[z]$) and thus $[x] = [y]$ as required. \square

Theorem 3.6.3. *Equality of terms $a, b \in \text{Tm}(\Gamma_{SK}, A)$ is undecidable.*

Proof. Suppose it were decidable; then for any $x, y \in \Lambda$ we can decide the equality of $\llbracket x \rrbracket, \llbracket y \rrbracket \in \text{Tm}(\Gamma_{SK}, A)$. By Lemma 3.6.1 and Theorem 3.6.2, $\llbracket x \rrbracket = \llbracket y \rrbracket$ if and only if $x \sim y$, so we can in turn decide the convertibility of SK-combinators, which is impossible. \square

3.6.2 The second proof: separating classes of Turing machines

In the first proof we reduce an undecidable problem to the judgmental equality of open terms, but establishing the completeness of this reduction requires appealing to the set-theoretic model of type theory. Our second proof relies only on the consistency of extensional type theory, showing that deciding judgmental equality of closed functions would allow us to algorithmically separate two recursively inseparable subsets of \mathbb{N} .

Notation 3.6.4. Fix a standard, effective Gödel encoding of Turing machines, in which the standard operations on Turing machines are definable by primitive recursion. We write ϕ_n for the partial function induced by the Turing machine encoded by n .

Theorem 3.6.5 (Rosser [Ros36], Trahtenbrot [Tra53], and Kleene [Kle50]). *Consider the following two subsets of the natural numbers:*

$$A = \{n \in \mathbb{N} \mid \phi_n(n) \text{ terminates with result } 0\}$$

$$B = \{n \in \mathbb{N} \mid \phi_n(n) \text{ terminates with result } 1\}$$

There is no Turing machine which terminates on all inputs and separates A from B .

Proof. Suppose we are given a Turing machine e which always terminates with value 0 or 1, such that $e(n) = 0$ when $n \in A$ and $e(n) = 1$ when $n \in B$. Consider the algorithm

$$F(n) := \begin{cases} \text{halt}(1) & e(n) = 0 \\ \text{halt}(0) & e(n) = 1 \end{cases}$$

Because e terminates on all inputs, so does F . Note that $e(F(n)) \neq e(n)$ by construction: if $e(F(n)) = 1$ then $e(n) = 0$ and vice versa. By the second recursion theorem, there exists a Turing machine f realizing F applied to its own Gödel number. However, $e(f)$ can be neither 0 nor 1 as $e(f) = e(F(f))$ by definition, but $e(f) \neq e(F(f))$. \square

We will show that the existence of a normalization structure for extensional type theory contradicts the above theorem. First, we observe that we can write a “small-step interpreter” for Turing machines in type theory. Let us write `TM` and `State` for `Nat` to indicate that we are interpreting a natural number as a Turing machine or Turing machine state respectively, as encoded by ϕ . Then we can define the following functions in type theory by primitive recursion:

- `init : TM \rightarrow Nat \rightarrow State`
- `hasHalted : State \rightarrow $\sum_{b:\text{Bool}}$ if (Nat, Unit, b)`
- `step : State \rightarrow State`

Using these operations, we can run a Turing machine for an arbitrary but finite number of steps on any input, determine whether it has halted, and if so, extract the result. We can therefore define the following function:

```
-- returns true iff Turing machine n halts on n with result 1 in fewer than t steps
returnOne : TM  $\rightarrow$  Nat  $\rightarrow$  Bool
returnOne n t = go (init n n) t
  where
    go : State  $\rightarrow$  Nat  $\rightarrow$  Bool
    go s zero = false
    go s (suc n) =
      if fst (hasHalted s) then isOne (snd (hasHalted s)) else go (step s) n
```

Let $H_0 \in \mathbb{N}$ be the encoding of a Turing machine which immediately halts with result 0 regardless of its input. Then, writing \bar{m} for the element of $\text{Tm}(1, \text{Nat})$ corresponding to $m \in \mathbb{N}$, we will show that `returnOne(\bar{n})`, `returnOne(\bar{H}_0)` \in $\text{Tm}(1, \Pi(\text{Nat}, \text{Bool}))$ are equal (resp., unequal) when n is a Turing machine which halts with result 0 (resp., 1).

Lemma 3.6.6. *If $n \in \mathbb{N}$ is such that $\phi_n(n) = 0$, then*

$$1 \vdash \text{returnOne } \bar{n} = \text{returnOne } \bar{H}_0 : \Pi(\text{Nat}, \text{Bool}).$$

Proof. By the η rule for Π -types, it suffices to show

$$1, t : \text{Nat} \vdash \text{returnOne } \bar{n} t = \text{returnOne } \bar{H}_0 t : \text{Bool}$$

By equality reflection, this follows from:

$$1, t : \text{Nat} \vdash P_t : \text{Eq}(\text{Bool}, \text{returnOne } \bar{n} t, \text{returnOne } \bar{H}_0 t)$$

In Exercise 3.22 the reader will establish this by Nat elimination on t . Note that by $\phi_n(n) = 0$, there exists some number ℓ such that the Turing machine encoded by n halts in t steps on n with result 0. Thus we must in essence construct the following terms:

$$1, t : \text{Nat} \vdash P_0 : \text{Eq}(\text{Bool}, \text{returnOne } \bar{n} \text{ zero}, \text{returnOne } \bar{H}_0 \text{ zero})$$

$$1, t : \text{Nat} \vdash P_1 : \text{Eq}(\text{Bool}, \text{returnOne } \bar{n} (\text{suc zero}), \text{returnOne } \bar{H}_0 (\text{suc zero}))$$

\vdots

$$1, t : \text{Nat} \vdash P_{\ell+1} : \text{Eq}(\text{Bool}, \text{returnOne } \bar{n} (\text{suc}^{\ell+1} t), \text{returnOne } \bar{H}_0 (\text{suc}^{\ell+1} t))$$

In the above, we write $\text{suc}^{\ell+1}(t)$ for the $(\ell + 1)$ -fold application of $\text{suc}(-)$ to t . When $i \leq \ell$ it is straightforward to construct P_i , as both sides equal **false**. For $P_{\ell+1}$, we note that $\text{returnOne } m (\text{suc}^k t) = \text{false}$ when m encodes a machine which halts in fewer than k steps with a result other than 1, completing the proof. \square

Exercise 3.22. Fill in the gap in the above argument using the elimination principle for Nat .

The remaining condition is easier to show.

Lemma 3.6.7. *If $n \in \mathbb{N}$ is such that $\phi_n(n) = 1$, then if the equality*

$$1 \vdash \text{returnOne } \bar{n} = \text{returnOne } \bar{H}_0 : \Pi(\text{Nat}, \text{Bool})$$

holds, extensional type theory is inconsistent.

Proof. Because $\phi_n(n)$ terminates, there is some number of steps t for which $\text{returnOne } \bar{n} t = \text{true}$. On the other hand, $\text{returnOne } \bar{H}_0 t = \text{false}$ for every t , so by applying both of these equal functions to t we conclude that $1 \vdash \text{true} = \text{false} : \text{Bool}$. By Theorem 2.6.3 this implies extensional type theory is inconsistent. \square

Theorem 3.6.8. *The judgmental equality $1 \vdash \text{returnOne } \bar{n} = \text{returnOne } \bar{H}_0 : \Pi(\text{Nat}, \text{Bool})$ cannot be decidable for all $n \in \mathbb{N}$.*

Proof. By Lemma 3.6.6, this equation holds if $\phi_n(n) = 0$; by Lemma 3.6.7 and Theorem 3.4.8, it does not hold if $\phi_n(n) = 1$. If this equation were decidable, we would be able to define a terminating algorithm which separates the subsets of $n \in \mathbb{N}$ for which $\phi_n(n) = 0$ and $\phi_n(n) = 1$, contradicting Theorem 3.6.5. \square

Further reading

There are a number of excellent pedagogical resources on type-checkers for dependent type theory that we encourage our implementation-inclined readers to explore. Coquand [Coq96] describes algorithms for bidirectional type-checking and deciding equality along with a proof sketch of correctness. Löh, McBride, and Swierstra [LMS10] include additional exposition and a complete Haskell implementation that extends a type-checker for a simply-typed calculus that is also described in the paper. The Mini-TT tutorial by Coquand et al. [Coq+09] includes a Haskell implementation of a type theory which is unsound (allowing arbitrary fixed-points) but supports data type declarations and basic pattern matching.

In addition to the aforementioned papers, there are numerous online resources, including a tutorial by Christiansen [Chr19] on the *normalization by evaluation* algorithm for deciding equality, and the *elaboration-zoo* of Kovács [Kov] which is an excellent resource for more advanced implementation techniques.

Intensional type theory

In Chapter 3 we outlined several key properties of type theories: consistency states that type theory can be viewed as a logic, canonicity states that type theory can be viewed as a programming language, normalization allows us to define a type-checking algorithm, and invertibility of type constructors improves that algorithm. Unfortunately, we also saw in Section 3.6 that extensional type theory does not satisfy the latter two properties due to the *equality reflection* rule of its **Eq**-types (Section 2.4.4).

If we remove **Eq**-types from extensional type theory then it will satisfy all four metatheorems above, but it becomes unusably weak. A foreseeable consequence is that type theory would no longer have an equality proposition; a more subtle issue is that many equations stop holding altogether, judgmentally or otherwise. This is because inductive types are characterized by maps into other *types* only, so what properties they enjoy depends on what types exist. Indeed we have already seen that **Eq**-types allow us to prove their η -rules and universes allow us to prove disjointness of their constructors; without **Eq**-types their η -rules will no longer be provable, and disjointness cannot even be stated!

We are left asking: *how should we internalize judgmental equality as a type, if not **Eq**?* This question has preoccupied type theorists for decades and—fortunately for their continued employment—has no clear-cut answer. We will find that deleting equality reflection causes equality types to become underconstrained, and their most canonical replacement, *intensional identity types*, lack several important reasoning principles. The decades-long quest for a suitable identity type has resulted in many subtle variations as well as some major innovations in type theory, as we will explore in Chapter 5. But first we turn our attention to *intensional type theory*, or type theory with intensional identity types, the system on which most type-theoretic proof assistants are based.

Notation 4.0.1. We adopt the common acronyms ETT and ITT for extensional type theory and intensional type theory respectively.

In this chapter In Section 4.1 we explore the basic properties that any propositional equality connective must satisfy, and show that a small set of primitive operations suffice to recover many of the positive consequences of equality reflection while allowing for normalization. In Section 4.2 we formally define the intensional identity type according to the framework of inductive types outlined in Section 2.5, and show that this type precisely satisfies the properties of equality outlined above. In Section 4.3 we compare extensional and intensional identity types, noting that the latter lacks several important principles, but by adding two axioms to it we can recover all the reasoning principles of extensional type theory in a precise sense. Finally, in Section 4.4, we summarize a line of research on

observational type theory [AMS07], which attempts to improve intensional identity types without sacrificing normalization.

Goals of the chapter By the end of this chapter, you will be able to:

- Define `subst` and contractibility of singletons, use them to prove other properties of equality, and implement them using intensional identity types.
- Explain how intensional identity types fit into the framework of internalizing judgmental structure that we developed in Chapter 2.
- Discuss the relationship and tradeoffs between intensional and extensional equality.
- Informally describe observational type theory, and explain how it addresses the shortcomings of intensional and extensional type theory.

4.1 Programming with propositional equality

In this section we will informally consider what properties should be satisfied by any “type of equations.” Recall from Section 1.1.3 that such a *propositional* (or *typal*, or *internal*) notion of equality is important for proving equations between types that type-checkers cannot handle automatically, and that such type equations allow us to cast (coerce) between the types involved. In Section 3.1 we discussed how type-checkers automatically handle definitional (judgmental) type equalities; one can therefore think of propositional type equalities as “verified casts” that users manually insert into terms.

Our starting point will be the type theory described in Chapter 2 but without `Eq`-types. Instead we will add an *identity type* `Id`¹ with the same formation (and universe introduction) rule but no other properties yet:

$$\frac{\Gamma \vdash a : A \quad \Gamma \vdash b : A}{\Gamma \vdash \mathbf{Id}(A, a, b) \text{ type}} \qquad \frac{\Gamma \vdash a : \mathbf{U}_i \quad \Gamma \vdash x : \mathbf{El}(a) \quad \Gamma \vdash y : \mathbf{El}(a)}{\Gamma \vdash \mathbf{id}(a, x, y) : \mathbf{U}_i \quad \Gamma \vdash \mathbf{El}(\mathbf{id}(a, x, y)) = \mathbf{Id}(\mathbf{El}(a), x, y) \text{ type}}$$

The primary way to use a proof of `Id(A, a, a')` is in concert with an A -indexed family of types $b : A \rightarrow \mathbf{U}$; namely, we conclude that the a and a' instances of this family are themselves equal in the sense that we have a proof of `Id(U, b a, b a')`, and as a result we are able to cast between the types `El(b a)` and `El(b a')`. Notably, because type equality is central to this story, universes will play a major role in this section.

¹Although beyond the scope of this book, we expect the **Superego** connective to internalize the rules of type theory; arguably singleton types internalize the self and thus serve as the **Ego**.

Notation 4.1.1. What should we call terms of type $\mathbf{Id}(A, a, b)$? This type will no longer precisely internalize the equality judgment so it can be misleading to call them *equalities* between a and b . On the other hand, calling them “*proofs of equality* between a and b ” is too cumbersome. We will refer to them as *identifications* between a and b .

Notation 4.1.2. In the remainder of this section we will return to the informal notation of Chapter 1; in particular, we omit $\mathbf{El}(-)$, thereby suppressing the difference between types and terms of type \mathbf{U} . We resume our more rigorous notation in Section 4.2.

4.1.1 Constructing identifications

Following the discussion above, we can already formulate two necessary conditions on $\mathbf{Id}(A, a, b)$. First, we must have some source of identifications between terms. As with \mathbf{Eq} -types we choose reflexivity; in concert with definitional equality, this allows us to prove any terms are identified as long as they differ only by β , η , and expanding definitions:

$$\mathbf{refl} : \{A : \mathbf{U}\} \rightarrow (a : A) \rightarrow \mathbf{Id}(A, a, a)$$

Secondly, given an identification $\mathbf{Id}(A, a, a')$ and a dependent type $B : A \rightarrow \mathbf{U}$, we must be able to convert terms of type $B(a)$ to $B(a')$, a process (confusingly) known as *substitution*:

$$\mathbf{subst} : \{A : \mathbf{U}\} \{a a' : A\} \rightarrow (B : A \rightarrow \mathbf{U}) \rightarrow \mathbf{Id}(A, a, a') \rightarrow B a \rightarrow B a'$$

Remark 4.1.3. The \mathbf{subst} function did not emerge in our discussion of \mathbf{Eq} -types for the simple reason that equality reflection trivializes it: $\mathbf{subst} B p b = b$. Indeed, all of the operations we discuss in this section are trivial in the presence of equality reflection. \diamond

By assuming that \mathbf{Id} -types satisfy \mathbf{refl} and \mathbf{subst} we are off to a good start, but *a priori* these are only two of the many combinators that we expect to be definable for $\mathbf{Id}(A, a, b)$; for starters, as an equality relation, identifications ought to be not only reflexive but also symmetric and transitive. Fortunately and somewhat surprisingly, it turns out that both symmetry and transitivity are consequences of \mathbf{refl} and \mathbf{subst} .

Lemma 4.1.4. *Using \mathbf{refl} and \mathbf{subst} , we can prove symmetry of identifications, i.e.,*

$$\mathbf{sym} : \{A : \mathbf{U}\} \{a b : A\} \rightarrow \mathbf{Id}(A, a, b) \rightarrow \mathbf{Id}(A, b, a)$$

Proof. Fix $A : \mathbf{U}$ and $a b : A$ and $p : \mathbf{Id}(A, a, b)$. To construct a term of type $\mathbf{Id}(A, b, a)$, we simply choose a clever B at which to instantiate \mathbf{subst} :

$$\begin{aligned} B &: A \rightarrow \mathbf{U} \\ B x &= \mathbf{id}(A, x, a) \end{aligned}$$

In particular, note that $B(a) = \mathbf{Id}(A, a, a)$ is easily proven by `refl`, and $B(b) = \mathbf{Id}(A, b, a)$ is our goal; thus `subst B p` is a function $B(a) \rightarrow B(b)$ and our goal follows soon after:

$$\begin{aligned} \text{sym} &: \{A : \mathbf{U}\} \{a\ b : A\} \rightarrow \mathbf{Id}(A, a, b) \rightarrow \mathbf{Id}(A, b, a) \\ \text{sym } \{A\ a\ b\} p &= \text{subst } (\lambda x \rightarrow \mathbf{id}(A, x, a))\ p\ (\text{refl } a) \end{aligned} \quad \square$$

Lemma 4.1.5. *Using `refl` and `subst`, we can prove transitivity of identifications, i.e.,*

$$\text{trans} : \{A : \mathbf{U}\} \{a\ b\ c : A\} \rightarrow \mathbf{Id}(A, a, b) \rightarrow \mathbf{Id}(A, b, c) \rightarrow \mathbf{Id}(A, a, c)$$

Proof. Fix $A : \mathbf{U}$, $a\ b\ c : A$, $p : \mathbf{Id}(A, a, b)$, and $q : \mathbf{Id}(A, b, c)$. To construct a term of type $\mathbf{Id}(A, a, c)$, we again choose a clever instantiation of `subst`, in this case $B(x) = \mathbf{Id}(A, a, x)$. Once again, $B(b)$ is easily proven by our assumption p , and $B(c)$ is our goal. Substituting along $q : \mathbf{Id}(A, b, c)$ completes our proof:

$$\begin{aligned} \text{trans} &: \{A : \mathbf{U}\} \{a\ b\ c : A\} \rightarrow \mathbf{Id}(A, a, b) \rightarrow \mathbf{Id}(A, b, c) \rightarrow \mathbf{Id}(A, a, c) \\ \text{trans } \{A\ a\ b\ c\} p\ q &= \text{subst } (\lambda x \rightarrow \mathbf{id}(A, a, x))\ q\ p \end{aligned} \quad \square$$

Exercise 4.1. Provide an alternative proof `trans'` of Lemma 4.1.5 which substitutes along p rather than q , using a slightly different choice of B .

In fact, `refl` and `subst` also allow us to prove that identifications are a congruence, in the sense that given $\mathbf{Id}(A, a, a')$ and $f : A \rightarrow B$, we obtain an identification $\mathbf{Id}(B, f\ a, f\ a')$.

Lemma 4.1.6. *Using `subst`, we can prove congruence of identifications, i.e.,*

$$\text{cong} : \{A\ B : \mathbf{U}\} \{a\ a' : A\} \rightarrow (f : A \rightarrow B) \rightarrow \mathbf{Id}(A, a, a') \rightarrow \mathbf{Id}(B, f\ a, f\ a')$$

Proof. The proof strategy remains the same, so we proceed directly to the term:

$$\begin{aligned} \text{cong} &: \{A\ B : \mathbf{U}\} \{a\ a' : A\} \rightarrow (f : A \rightarrow B) \rightarrow \mathbf{Id}(A, a, a') \rightarrow \mathbf{Id}(B, f\ a, f\ a') \\ \text{cong } \{A\ B\ a\ a'\} f\ p &= \text{subst } (\lambda x \rightarrow \mathbf{id}(B, f\ a, f\ x))\ p\ (\text{refl } (f\ a)) \end{aligned} \quad \square$$

Finally, we must consider how `subst` ought to compute. Because `subst` can produce terms of any type, including `Bool` and `Nat`, we must impose some definitional equalities on it if our type theory is to satisfy canonicity (Section 3.4). One equation springs to mind immediately: if we apply `subst B` to `refl a`, the resulting coercion $B\ a \rightarrow B\ a$ has the type of the identity function, so it is reasonable to ask for it to *be* the identity function. That is, we ask for the following definitional equality:

$$\text{subst } B\ (\text{refl } a)\ b = b : B\ a$$

4.1.2 Constructing identifications of identifications

Although `refl` and `subst` go quite a long way, they *do not* suffice to derive all the properties of identifications we might expect; we start encountering their limits as soon as we consider identifications between elements of $\mathbf{Id}(A, a, b)$ itself. These *identifications of identifications* arise very naturally in practice. Quite often we must use `subst` when constructing a dependently-typed term in order to align various type indices; if we ever construct a type that depends on such a term, we will very quickly be in the business of proving that two potentially distinct sequences of `subst` casts are themselves equal.

For the sake of concreteness, consider the following pair of operations that “rotate” a `Vector` (a list of specified length, as defined in Chapter 1):

$$\begin{aligned} \text{append} &: \{A : \mathbf{U}\} \{n\ m : \mathbf{Nat}\} \rightarrow \mathbf{Vec}\ A\ n \rightarrow \mathbf{Vec}\ A\ m \rightarrow \mathbf{Vec}\ A\ (n + m) \\ \text{comm} &: \{n\ m : \mathbf{Nat}\} \rightarrow \mathbf{Id}(\mathbf{Nat}, n + m, m + n) \end{aligned}$$

$$\begin{aligned} \text{rot1} &: \{A : \mathbf{U}\} \{n : \mathbf{Nat}\} \rightarrow \mathbf{Vec}\ A\ n \rightarrow \mathbf{Vec}\ A\ n \\ \text{rot1}\ [] &= [] \\ \text{rot1}\ \{A\ n\}\ (x :: xs) &= \text{subst}\ (\mathbf{Vec}\ A)\ (\text{comm}\ n\ 1)\ (\text{append}\ xs\ (x :: [])) \end{aligned}$$

$$\begin{aligned} \text{rot2} &: \{A : \mathbf{U}\} \{n : \mathbf{Nat}\} \rightarrow \mathbf{Vec}\ A\ n \rightarrow \mathbf{Vec}\ A\ n \\ \text{rot2}\ [] &= [] \\ \text{rot2}\ (x :: []) &= x :: [] \\ \text{rot2}\ \{A\ n\}\ (x_0 :: x_1 :: xs) &= \text{subst}\ (\mathbf{Vec}\ A)\ (\text{comm}\ n\ 2)\ (\text{append}\ xs\ (x_0 :: x_1 :: [])) \end{aligned}$$

We expect to be able to prove that `rot1` twice is the same as `rot2`:

$$\{A : \mathbf{U}\} \{n : \mathbf{Nat}\} \rightarrow (xs : \mathbf{Vec}\ A\ (2 + n)) \rightarrow \mathbf{Id}(\mathbf{Vec}\ A\ (2 + n), \text{rot1}\ (\text{rot1}\ xs), \text{rot2}\ xs)$$

However, this will not be possible with our current set of primitives. In our definitions of `rot1` and `rot2` we were forced to include various applications of `subst` to correct mismatches between the indices $(n + 1)$, $(1 + n)$ and $(n + 2)$, $(2 + n)$, and these `subst` terms will get in our way as we try to establish the above identification. If we proceed by induction on `xs`, for instance, we will get stuck attempting to construct an identification between

$$\begin{aligned} &\text{subst}\ (\mathbf{Vec}\ A)\ (\text{comm}\ n\ 1) \\ &\quad (\text{append}\ (\text{subst}\ (\mathbf{Vec}\ A)\ (\text{comm}\ n\ 1)\ (\text{append}\ xs\ (x_0 :: [])))\ (x_1 :: [])) \end{aligned}$$

and

$$\text{subst}\ (\mathbf{Vec}\ A)\ (\text{comm}\ n\ 2)\ (\text{append}\ xs\ (x_0 :: x_1 :: []))$$

of type $\mathbf{Vec}\ A\ (2+n)$. Unfortunately, because `n` is a variable, neither `comm n 1` nor `comm n 2` are the reflexive identification, so we can make no further progress.

The above example is a bit involved, but there are many smaller (albeit more contrived) examples of identifications that are beyond our reach; for example, given a variable $p : \mathbf{Id}(A, a, b)$ we cannot construct an identification $\mathbf{Id}(\mathbf{Id}(A, a, b), p, \mathbf{sym}(\mathbf{sym} p))$.

Our “API” for identity types is thus missing an operation that allows us to prove identifications between two identifications. To hit upon this operation, we introduce the concept of (*propositional*) *singleton types* (in contrast to the “definitional singleton types” of Section 3.3). Given a type A and a term $a : A$, the singleton type $[a]$ is defined as follows:

$$[a] = \sum_{b:A} \mathbf{Id}(A, a, b)$$

That is, $[a]$ is the type of “elements of A that can be identified with a .” Intuitively, there should only be one such element, namely a itself—or to be more precise, $(a, \mathbf{refl} a)$. But this, too, is not yet provable. Certainly, given an arbitrary element $(b, p) : [a]$ we can see that (by p) their first projections a and b are identified, but we have no way of identifying their second projections $\mathbf{refl} a$ and p .

In fact, most of our “coherence problems” of identifying identifications can be reduced to the problem of identifying all elements of $[a]$: this is in some sense the *ur*-coherence problem. Intuitively this is because being able to identify arbitrary (b, p) with $(a, \mathbf{refl} a)$ allows us to transform subst terms involving the arbitrary identification p into subst terms involving the distinguished identification $\mathbf{refl} a$, the latter of which “compute away.”

Lemma 4.1.7. *Suppose we are given some $A : \mathbf{U}$ and $a : A$ such that all elements of $[a]$ are identified; then for any $b : A$ and $p : \mathbf{Id}(A, a, b)$ we have $\mathbf{Id}(\mathbf{Id}(A, a, b), p, \mathbf{sym}(\mathbf{sym} p))$.*

Proof. Fixing A, a, b , and p , we notice that $(a, \mathbf{refl} a), (b, p) : [a]$ by definition, and thus by assumption we have an identification $q : \mathbf{Id}([a], (a, \mathbf{refl} a), (b, p))$. As before, we shall choose a clever B for which subst B solves our problem, namely:

$$\begin{aligned} B &: [a] \rightarrow \mathbf{U} \\ B(b_0, p_0) &= \mathbf{id}(\mathbf{id}(A, a, b_0), p_0, \mathbf{sym}(\mathbf{sym} p_0)) \end{aligned}$$

Inspecting our definition of Lemma 4.1.4, we see that $\mathbf{sym}(\mathbf{refl} x) = \mathbf{refl} x$ definitionally, and thus the following definitional equalities hold:

$$\begin{aligned} B(a, \mathbf{refl} a) &= \mathbf{Id}(\mathbf{Id}(A, a, a), \mathbf{refl} a, \mathbf{sym}(\mathbf{sym}(\mathbf{refl} a))) \\ &= \mathbf{Id}(\mathbf{Id}(A, a, a), \mathbf{refl} a, \mathbf{sym}(\mathbf{refl} a)) \\ &= \mathbf{Id}(\mathbf{Id}(A, a, a), \mathbf{refl} a, \mathbf{refl} a) \\ B(b, p) &= \mathbf{Id}(\mathbf{Id}(A, a, b), p, \mathbf{sym}(\mathbf{sym} p)) \end{aligned}$$

It is easy to produce an element of the former type (namely, $\mathbf{refl}(\mathbf{refl} a)$), the latter type is our goal, and q is an identification between the two indices. Thus:

$$\begin{aligned}
& \text{symsym} : \{A : \mathbf{U}\} \{a \ b : A\} \rightarrow (p : \mathbf{Id}(A, a, b)) \rightarrow \mathbf{Id}(\mathbf{Id}(A, a, b), p, \text{sym}(\text{sym } p)) \\
& \text{symsym } \{A \ a \ b\} \ p = \text{subst} \\
& \quad (\lambda(b_0, p_0) \rightarrow \mathbf{id}(\mathbf{id}(A, a, b_0), p_0, \text{sym}(\text{sym } p_0))) \\
& \quad ? : \mathbf{Id}([a], (a, \text{refl } a), (b, p)) \quad \text{-- by assumption} \\
& \quad (\text{refl}(\text{refl } a)) \quad \square
\end{aligned}$$

We substantiate the assumption of Lemma 4.1.7 with a new primitive operation on identity types, `uniq`, that identifies $(a, \text{refl } a)$ with arbitrary elements of $[a]$. (By `sym` and `trans`, it follows that any two arbitrary elements of $[a]$ are also identified.) As with `subst`, we also assert that a certain definitional equality holds when `uniq` is supplied with the reflexive identification. This operation is often called *singleton contractibility* [Coq14; UF13], and it will feature prominently in Chapter 5.

$$\begin{aligned}
& \text{uniq} : \{A : \mathbf{U}\} \{a : A\} \rightarrow (x : [a]) \rightarrow \mathbf{Id}([a], (a, \text{refl } a), x) \\
& \quad \text{uniq } (a, \text{refl } a) = \text{refl } (a, \text{refl } a)
\end{aligned}$$

Exercise 4.2. Like `subst`, `uniq` is definable in extensional type theory; show this.

Exercise 4.3. Recalling `trans` (Lemma 4.1.5) and `trans'` (Exercise 4.1), use `subst` and `uniq` to construct a term of the following type:

$$\begin{aligned}
& \{A : \mathbf{U}\} \{a \ b \ c : A\} \rightarrow (p : \mathbf{Id}(A, a, b)) \rightarrow (q : \mathbf{Id}(A, b, c)) \rightarrow \\
& \quad \mathbf{Id}(\mathbf{Id}(A, a, c), \text{trans } p \ q, \text{trans}' \ p \ q)
\end{aligned}$$

4.1.3 Intensional identity types

To summarize Sections 4.1.1 and 4.1.2, we have asked for $\mathbf{Id}(A, a, b)$ to support the following three operations subject to two definitional equalities:

$$\begin{aligned}
& \text{refl} : \{A : \mathbf{U}\} \rightarrow (a : A) \rightarrow \mathbf{Id}(A, a, a) \\
& \text{subst} : \{A : \mathbf{U}\} \{a \ a' : A\} \rightarrow (B : A \rightarrow \mathbf{U}) \rightarrow \mathbf{Id}(A, a, a') \rightarrow B \ a \rightarrow B \ a' \\
& \text{uniq} : \{A : \mathbf{U}\} \{a : A\} \rightarrow (x : [a]) \rightarrow \mathbf{Id}([a], (a, \text{refl } a), x) \\
& \\
& \text{subst } B \ (\text{refl } a) \ b = b \\
& \text{uniq}(a, \text{refl } a) = \text{refl } (a, \text{refl } a)
\end{aligned}$$

Definition 4.1.8. An *intensional identity type* is any type $\mathbf{Id}(A, a, b)$ equipped with the three operations above satisfying the two definitional equalities above.

Intensional identity types were introduced by Martin-Löf [ML75] and have been the “standard” formulation of propositional equality in type theory for most of the intervening years, although various authors have presented them via different but equivalent sets of primitive operations and equations [CP90b; PPM90; PM93; Str93; Coq14].² Our presentation most closely follows Coquand [Coq14] which, to our knowledge, was first proposed by Steve Awodey in 2009. In Sections 4.3 and 4.4 we will also consider related but *non-equivalent* presentations endowing $\mathbf{Id}(A, a, b)$ with more properties [Str93; Hof95a; AMS07].

Let us be clear, however, that this broad agreement in the literature is not an indication of happiness. On the contrary, most type theorists have many complaints about intensional identity types: there are several important properties that they do *not* satisfy, and they can be frustrating in practice for a number of reasons. They have persisted for so long because of a relative lack of compelling alternatives that also satisfy the two crucial properties of:

1. Capturing the most important properties of equality—reflexivity, symmetry, transitivity, congruence, substitutivity, etc.—thus enabling a wide range of constructions.
2. Their inclusion in a type theory is compatible with all the metatheorems discussed in Chapter 3, especially—unlike \mathbf{Eq} -types—normalization.

In Section 4.3 we will discuss the shortcomings of \mathbf{Id} -types in more detail, but it will turn out that these shortcomings can be mostly overcome by adding several axioms (postulated terms, or in essence, free variables) to type theory. Adding such axioms causes canonicity to fail, but as discussed in Section 3.4, type theories without canonicity are merely frustrating (requiring more manual reasoning by identifications), whereas type theories without normalization are essentially un-type-checkable. As a result, many users of type theory opt to work with \mathbf{Id} -types with some additional axioms.

But before we get ahead of ourselves, we proceed by formally defining \mathbf{Id} -types and thus the type theory known as *intensional type theory*.

4.2 Intensional identity types

In this section we formally define intensional identity types, or \mathbf{Id} -types, returning to the style of definition adopted throughout Chapter 2. Although it is possible to add \mathbf{Id} -types to extensional type theory, we are primarily interested in defining *intensional type theory*, which is obtained by replacing certain rules of ETT by the rules in this section. Specifically, we remove from the theory of Chapter 2 all rules pertaining to \mathbf{Eq} -types; in Appendix A those rules are annotated (ETT), and the rules added in this section are annotated (ITT).

²The equivalence between the presentations of Martin-Löf [ML75] and Paulin-Mohring [PM93] is due to Hofmann [Str93, Addendum].

Although the rules for **Id**-types appear complicated and unmotivated at first, it will turn out that they arise naturally from our methodology that types internalize judgmental structure. Recalling Slogan 2.5.3, connectives in type theory are specified by a natural type-forming operation whose terms are either defined by a mapping-in property (a natural isomorphism with judgmentally-defined structure) or a mapping-out property (an algebra signature for which the type carries a weakly initial algebra).

The formation rule of $\mathbf{Id}(A, a, b)$ is identical to that of $\mathbf{Eq}(A, a, b)$:

$$\frac{\Gamma \vdash a : A \quad \Gamma \vdash b : A}{\Gamma \vdash \mathbf{Id}(A, a, b) \text{ type}} \quad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : A \quad \Gamma \vdash b : A}{\Delta \vdash \mathbf{Id}(A, a, b)[\gamma] = \mathbf{Id}(A[\gamma], a[\gamma], b[\gamma]) \text{ type}}$$

Or equivalently, the following type-forming operation natural in Γ :

$$\mathbf{Id}_\Gamma : (\sum_{A \in \mathbf{Ty}(\Gamma)} \mathbf{Tm}(\Gamma, A) \times \mathbf{Tm}(\Gamma, A)) \rightarrow \mathbf{Ty}(\Gamma)$$

We must now decide whether to define $\mathbf{Id}(A, a, b)$ by a mapping-in property or a mapping-out property. In Chapter 2 we saw that mapping-in properties are generally both simpler and better-behaved, but we already defined **Eq**-types by the mapping-in property of internalizing judgmental equality (i.e., $\mathbf{Tm}(\Gamma, \mathbf{Eq}(A, a, b)) \cong \{\star \mid a = b\}$), and it is unclear what other structure we could ask for **Id**-types to internalize.³

Faced with no other options, we are forced to consider a mapping-out property instead. Per the discussion in Sections 2.5.2 and 2.5.3, such a property starts with a collection of natural term constructors of $\mathbf{Id}(A, a, b)$, in this case only **reflexivity**:

$$\frac{\Gamma \vdash a : A}{\Gamma \vdash \mathbf{refl} : \mathbf{Id}(A, a, a)} \quad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : A}{\Delta \vdash \mathbf{refl}[\gamma] = \mathbf{refl} : \mathbf{Id}(A[\gamma], a[\gamma], a[\gamma])}$$

Or equivalently, the following term-forming operation natural in Γ :

$$\mathbf{refl}_{\Gamma, A, a} \in \mathbf{Tm}(\Gamma, \mathbf{Id}(A, a, a))$$

Whereas the mapping-in property of **Eq**-types asserts that **refl** is their only inhabitant, the mapping-out property of **Id**-types will assert that every type *believes* that **refl** is their only inhabitant, in just the same way that every type “believes” that **true** and **false** are the only elements of **Bool**, namely that to map out of **Bool** it suffices to explain how to behave on **true** and **false**.

Remark 4.2.1. Like the induction principles of inductive types, the **subst** and **uniq** primitives of Section 4.1 are both maps out of $\mathbf{Id}(A, a, b)$ that have prescribed behavior on

³Cubical type theory in fact invents a new judgmental structure for propositional equality to internalize, but we will return to this point in Section 5.4.

the constructor **refl**. We will see shortly that both **subst** and **uniq** are definable via the **Id**-elimination principle we are about to present, and remarkably, that **Id**-elimination can conversely be recovered as a combination of **subst** and **uniq**!

Compared to **subst** and **uniq**, **Id**-elimination is more clearly motivated by general considerations (mapping-out properties), more self-contained (not requiring Σ -types), and even often more ergonomic in practice. But **subst** and **uniq** are nevertheless very important combinators that certainly merit special discussion. \diamond

Luckily **refl** is not a recursive constructor, so we can avoid the discussion of displayed algebras of Section 2.5.3 and return to the simpler characterization of mapping-out properties in Sections 2.5.2 and 2.5.4 as a section (right inverse) to substitution of constructors. In exchange we must for the first time consider an inductive type former that has formation data, namely a type A and two terms $a, b : A$.

Suppose we have a dependent type over an identity type:

$$\Gamma.A.A[\mathbf{p}].\mathbf{Id}(A[\mathbf{p}^2], \mathbf{q}[\mathbf{p}], \mathbf{q}) \vdash C \text{ type}$$

Into any term of the above type we can substitute **refl**:

$$(\mathbf{id.q.refl})^* : \mathbf{Tm}(\Gamma.A.A[\mathbf{p}].\mathbf{Id}(A[\mathbf{p}^2], \mathbf{q}[\mathbf{p}], \mathbf{q}), C) \rightarrow \mathbf{Tm}(\Gamma.A, C[\mathbf{id.q.refl}])$$

The elimination principle for **Id**-types is precisely a section of the above map.

Let us unpack this a bit. First, we rewrite the above map using named variables:

$$[a/b, \mathbf{refl}/p] : \mathbf{Tm}(\Gamma, a : A, b : A, p : \mathbf{Id}(A, a, b), C(a, b, p)) \rightarrow \mathbf{Tm}(\Gamma, a : A, C(a, a, \mathbf{refl}(a)))$$

A section to this map tells us that to construct an element of $C(a, b, p)$ for any $a, b : A$ and $p : \mathbf{Id}(A, a, b)$, it suffices to say what to do on a, a, \mathbf{refl} (i.e., provide a term of type $C(a, a, \mathbf{refl})$). Compared to our definition of **if** in Section 2.5.2, the context on the left is more complex because the domain of a dependent type $C : \mathbf{Id}(A, a, b) \rightarrow \mathbf{U}$ is itself dependent on $a, b : A$, and the context on the right is more complex because the constructor **refl** is dependent on $a : A$.

Remark 4.2.2. From a more nuts-and-bolts perspective, imagine that we asked for C not to be dependent on all three of a, b, p as $\Gamma, a : A, b : A, p : \mathbf{Id}(A, a, b) \vdash C(a, b, p)$ type, but only on p , i.e., $\Gamma, p : \mathbf{Id}(A, a, b) \vdash C(p)$ type for some fixed $a, b : A$. Then we would not even be able to even *state* what it means to substitute **refl** for p , because **refl** only has type $\mathbf{Id}(A, a, b)$ when a and b are definitionally equal. Instead, we ask for all of a, b, p to be variables, and consider the substitution of a, a, \mathbf{refl} for a, b, p . \diamond

Unfolding the above section into inference rules, we once again “build in a cut” by applying the stipulated term in context $\Gamma.A.A[\mathbf{p}].\mathbf{Id}(A[\mathbf{p}^2], \mathbf{q}[\mathbf{p}], \mathbf{q})$ to arguments $a : A$, $b : A$, and $p : \mathbf{Id}(A, a, b)$ all in context Γ . The first rule below is the section map itself, the

second rule is naturality of the section map, and the third states that applying the section map followed by $(\mathbf{id.q.refl})^*$ is the identity:

$$\frac{\Gamma \vdash a : A \quad \Gamma \vdash b : A \quad \Gamma \vdash p : \mathbf{Id}(A, a, b) \quad \Gamma.A.A[\mathbf{p}].\mathbf{Id}(A[\mathbf{p}^2], \mathbf{q}[\mathbf{p}], \mathbf{q}) \vdash C \text{ type} \quad \Gamma.A \vdash c : C[\mathbf{id.q.refl}]}{\Gamma \vdash \mathbf{J}(c, p) : C[\mathbf{id.a.b.p}]}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : A \quad \Gamma \vdash b : A \quad \Gamma \vdash p : \mathbf{Id}(A, a, b) \quad \Gamma.A.A[\mathbf{p}].\mathbf{Id}(A[\mathbf{p}^2], \mathbf{q}[\mathbf{p}], \mathbf{q}) \vdash C \text{ type} \quad \Gamma.A \vdash c : C[\mathbf{id.q.refl}]}{\Delta \vdash \mathbf{J}(c, p)[\gamma] = \mathbf{J}(c[(\gamma \circ \mathbf{p}).\mathbf{q}], p[\gamma]) : C[\gamma.a[\gamma].b[\gamma].p[\gamma]]}$$

$$\frac{\Gamma \vdash a : A \quad \Gamma.A.A[\mathbf{p}].\mathbf{Id}(A[\mathbf{p}^2], \mathbf{q}[\mathbf{p}], \mathbf{q}) \vdash C \text{ type} \quad \Gamma.A \vdash c : C[\mathbf{id.q.refl}]}{\Gamma \vdash \mathbf{J}(c, \mathbf{refl}) = c[\mathbf{id.a}] : C[\mathbf{id.a.a.refl}]}$$

These rules complete our definition of \mathbf{Id} -types and thus of intensional type theory.

As with the eliminators of \mathbf{Void} , \mathbf{Bool} , and \mathbf{Nat} , it can be helpful to think of $\mathbf{J}(c, p)$ as somehow “pattern-matching on p ” with clause c .

match (a, b, p) with
 $(a, a, \mathbf{refl}) \rightarrow c a$

From this perspective, the definitional equality $\mathbf{J}(c, \mathbf{refl}) = c[\mathbf{id.a}]$ states that the entire **match** expression reduces to c when (a, b, p) is indeed of the form (a, a, \mathbf{refl}) .

Remark 4.2.3. The name of \mathbf{J} for \mathbf{Id} -elimination dates back to Martin-Löf [ML84a], in which Martin-Löf notates \mathbf{Id} -types as \mathbf{I} , and he seems to have chosen \mathbf{J} simply because it is the next letter of the alphabet. At any rate, unlike **Identity** or **reflexivity**, it has no obvious meaning as the initial letter of pre-existing mathematical terminology.

For readers who might find this notational choice to be singularly arbitrary, we recall Scott’s story of mailing Church a postcard asking why λ was chosen as the symbol for function abstraction in his λ -calculus, and receiving the same postcard with the annotation “eeny, meeny, miny, moe” [Sco18]. \diamond

Like extensional type theory, intensional type theory satisfies consistency and canonicity; unlike extensional type theory, it also satisfies the metatheorems on open terms discussed in Chapter 3 and is therefore exceptionally well-behaved from the perspective of both theory and implementability.

Theorem 4.2.4 (Martin-Löf [ML71], Martin-Löf [ML75], and Coquand [Coq91]). *Intensional type theory satisfies consistency, canonicity, normalization, and has invertible type constructors.*

One typically deduces all of these properties from the proof of normalization: given that normalization amounts to concretely characterizing the sets $\text{Tm}(\Gamma, A)$ for all Γ, A , consistency and canonicity amount to verifying that these characterizations of $\text{Tm}(\mathbf{1}, \mathbf{Void})$ and $\text{Tm}(\mathbf{1}, \mathbf{Bool})$ contain zero and two elements respectively, and invertibility of Π -types amounts to inverting the induced $\Pi(-, -)$ map on normal forms. There are many proofs of normalization for intensional type theory and minor variations on it, some relying on semantic model constructions [AK16; Coq19; Ste21] and others more closely connected to algorithms used in real implementations [ACD07; Abe13; AÖV17].

From J to subst and uniq We close this section by showing that **J** is interprovable with the combination of `subst` and `uniq`, first that both `subst` and `uniq` are instances of **J**.

Notation 4.2.5. Our $\mathbf{J}(b, p)$ notation is not well-suited to informal constructions with named variables, because b silently binds a variable of type A , and moreover, the type C can be hard to infer by inspection. In our informal notation we will therefore wrap **J** as a function with the following type, satisfying the definitional equality $\mathbf{j} B b a a \mathbf{refl} = b a$.

$$\mathbf{j} : \{A : \mathbf{U}\} (C : (a b : A) \rightarrow \mathbf{Id}(A, a, b) \rightarrow \mathbf{U}) \rightarrow ((a : A) \rightarrow C a a \mathbf{refl}) \rightarrow (a b : A) (p : \mathbf{Id}(A, a, b)) \rightarrow C a b p$$

Likewise we introduce the functions $\mathbf{pi}, \mathbf{sig} : (A : \mathbf{U}) (B : A \rightarrow \mathbf{U}) \rightarrow \mathbf{U}$ as wrappers for the codes $\mathbf{pi}(-, -)$ and $\mathbf{sig}(-, -)$ respectively.

Exercise 4.4. Use the elimination principle **J** to define the function \mathbf{j} above, and check that your definition of \mathbf{j} satisfies the stipulated definitional equality.

The flexibility and complexity of **J** come from the fact that the *motive* [McB02] C can depend not only on the two elements of A but also the identification itself, both in arbitrary ways; many principles fall immediately out of **J** given a sufficiently clever choice of C .

Lemma 4.2.6. *Using \mathbf{j} we can define `subst`, i.e., a term of type*

$$\mathbf{subst} : \{A : \mathbf{U}\} \{a a' : A\} \rightarrow (B : A \rightarrow \mathbf{U}) \rightarrow \mathbf{Id}(A, a, a') \rightarrow B a \rightarrow B a'$$

satisfying the definitional equality $\mathbf{subst} \mathbf{refl} b = b$.

Proof. We will apply \mathbf{j} to the same $a, a' : A$ and $p : \mathbf{Id}(A, a, a')$ as `subst`, choosing a motive such that the type of the fully-applied \mathbf{j} will be $B a \rightarrow B a'$:

$$C x y _ = \mathbf{pi} (B x) (\lambda _ \rightarrow B y)$$

We have $C a a' p = B a \rightarrow B a'$ as desired, and it remains only to exhibit a term of type $(a : A) \rightarrow C a a \mathbf{refl} = (a : A) \rightarrow B a \rightarrow B a$, which is easy to do. In total:

$$\begin{aligned} \text{subst} &: \{A : \mathbf{U}\} \{a \ a' : A\} \rightarrow (B : A \rightarrow \mathbf{U}) \rightarrow \mathbf{Id}(A, a, a') \rightarrow B \ a \rightarrow B \ a' \\ \text{subst} &\{A \ a \ a'\} B \ p = j \ (\lambda x \ y \ _ \rightarrow \text{pi} (B \ x) (\lambda _ \rightarrow B \ y)) (\lambda _ \ x \rightarrow x) \ a \ a' \ p \end{aligned}$$

The reader can verify that the stipulated definitional equality holds. \square

Exercise 4.5. Check that the above definition of `subst` satisfies the required equation.

Lemma 4.2.7. *Using `j` we can define `uniq`, i.e., a term of type*

$$\text{uniq} : \{A : \mathbf{U}\} \{a : A\} \rightarrow (x : [a]) \rightarrow \mathbf{Id}([a], (a, \text{refl}), x)$$

satisfying the definitional equality `uniq (a, refl) = refl`.

Proof. Writing $A : \mathbf{U}$, $a : A$, and $x := (b, p) : \sum_{b:A} \mathbf{Id}(A, a, b)$ for the arguments of `uniq`, we will apply `j` to a, b, p with a motive that allows us to reduce the general case of a, b, p to the particular and easy case of a, a, refl :

$$C \ x \ y \ p' = \text{id}(\text{sig } A \ (\lambda z \rightarrow \text{id}(A, x, z)), (x, \text{refl}), (y, p'))$$

Then $C \ a \ b \ p = \mathbf{Id}([a], (a, \text{refl}), (b, p))$, and it remains only to exhibit a term of type $(a : A) \rightarrow C \ a \ a \ \text{refl} = (a : A) \rightarrow \mathbf{Id}([a], (a, \text{refl}), (a, \text{refl}))$, which is again easy:

$$\begin{aligned} \text{uniq} &: \{A : \mathbf{U}\} \{a : A\} \rightarrow (x : [a]) \rightarrow \mathbf{Id}([a], (a, \text{refl}), x) \\ \text{uniq} &\{A \ a\} (b, p) = j \ (\lambda x \ y \ p' \rightarrow \text{id}(\text{sig } A \ (\lambda z \rightarrow \text{id}(A, x, z)), (x, \text{refl}), (y, p'))) \\ &\quad (\lambda x \rightarrow \text{refl}_{(x, \text{refl})}) \ a \ b \ p \end{aligned}$$

The reader can again verify that the stipulated definitional equality holds.

Note that unlike the motive we used in Lemma 4.2.6, the motive here depends not only on $x, y : A$ but also the identification $p' : \mathbf{Id}(A, x, y)$. Note also that the motive actually generalizes our goal: rather than proving that for a fixed $a : A$ we can identify (a, refl) and $(b, p) : [a]$, we prove that for any $x, y : A$ we can identify (x, refl) and $(y, p') : [x]$. \square

Exercise 4.6. Check that the above definition of `uniq` satisfies the required equation.

And back again Conversely, using `subst` and `uniq` it is also possible to define a term `j` satisfying the required definitional equality. We leave most of the construction to the reader in the following series of exercises. In these exercises we fix the arguments of `j` as $A : \mathbf{U}$, $C : (a \ b : A) (p : \mathbf{Id}(A, a, b)) \rightarrow \mathbf{U}$, $c : (a : A) \rightarrow (C \ a \ a \ \text{refl})$, $a, b : A$, and $p : \mathbf{Id}(A, a, b)$, and we define the following “partially uncurried” type family:

$$\begin{aligned} C_a &: (x : [a]) \rightarrow \mathbf{U} \\ C_a \ x &= C \ a \ (\text{fst } x) \ (\text{snd } x) \end{aligned}$$

Exercise 4.7. Define a term $c_a : C_a (a, \text{refl})$.

Exercise 4.8. Without using **J**, define a term $q : \text{Id}([a], (a, \text{refl}), (b, p))$.

Exercise 4.9. Using c_a and q but not **J**, define a term $c_b : C_a (b, p)$.

Exercise 4.10. Show that the type of c_b is equal to $C a b p$, and use this to combine the previous three exercises into a definition of j that uses `subst` and `uniq` but not **J**.

Exercise 4.11. Check that your solution to Exercise 4.10 satisfies $j C c a a \text{refl} = c a$.

Exercise 4.12. We have seen in Remark 4.1.3 and Exercise 4.2 that `subst` and `uniq` are definable for **Eq**-types in ETT; from Exercise 4.10 it follows that j is also definable in ETT for **Eq**-types. Give an explicit definition of j for **Eq**-types in ETT. (Hint: you can combine the above results, but it is also fairly straightforward to arrive at the definition independently.)

Although it is perhaps easier to wrap one's head around `subst` and `uniq` rather than **J**, as we noted in Remark 4.2.1 it is often more straightforward in practice to use **J** directly. Consider for instance the function `cong` from Section 4.1, which we really ought to have stated for *dependent* functions:

$$\text{cong} : (f : (a : A) \rightarrow B a) \{a_0 a_1 : A\} (p : \text{Id}(A, a_0, a_1)) \rightarrow \\ \text{Id}(B a_1, \text{subst } B p (f a_0), f a_1)$$

Defining dependent `cong` in terms of `subst` and `uniq` is a headache, because one must use both simultaneously to handle the occurrence of p in the type. It is, however, straightforward to define with **J**:

$$\text{cong } f = j (\lambda a_0 a_1 p \rightarrow \text{Id}(B a_1, \text{subst } B p (f a_0), f a_1)) (\lambda a \rightarrow \text{refl}_{f(a)})$$

4.3 Limitations of the intensional identity type

We have now seen that the rules for **Id**-types are well-motivated from a theoretical perspective as the mapping-out formulation of equality, and that they support the operations of `subst` and `uniq` presented in Section 4.1, which in turn imply many properties including the symmetry, transitivity, and congruence of equality. We have also seen that ITT is more well-behaved than ETT (Theorem 4.2.4), and that all the rules of **Id**-types are validated by the **Eq**-types of ETT (Exercise 4.12).

Have we even lost anything at all by moving from ETT to ITT? Well, yes; the entire point of moving to ITT was to remove equality reflection from our theory, in light of its

undecidability (Section 3.6). Removing equality reflection does come at a cost: in ETT whenever we can prove $p : \mathbf{Eq}(A, a, a')$ we can freely use terms of type $B a$ at type $B a'$, but in ITT we must explicitly appeal to the proof p with $\text{subst } B p : B a \rightarrow B a'$.

So then are types and terms of ITT simply more *bureaucratic* than those of ETT, or does ITT actually “prove fewer statements” than ETT in some meaningful sense? This is an excellent question, and one that requires some care to set up precisely.

Given that closed types (of a consistent type theory) can be seen as logical propositions and their terms as their proofs, we might naïvely wonder *is every non-empty closed type of ETT also non-empty in ITT?* This question does not make sense as posed because, by equality reflection, well-formed types in ETT need not be well-formed in ITT. Consider for instance the following closed type of ETT:

$$(p : \mathbf{Eq}(\mathbf{Bool}, \mathbf{true}, \mathbf{false})) \rightarrow \mathbf{Eq}(\mathbf{Eq}(\mathbf{Bool}, \mathbf{true}, \mathbf{false}), \mathbf{refl}, p)$$

On the other hand, closed types of ITT *do* correspond to closed types of ETT in a more-or-less straightforward way, because their rules differ only in their choice of equality type, and the \mathbf{Eq} -types of ETT satisfy all the rules of the \mathbf{Id} -types of ITT (and more); to make this translation precise we once again turn to model theory.

Definition 4.3.1. We define a *model of ITT*, a *homomorphism of models of ITT*, and the *syntactic model \mathcal{T}_{ITT} of ITT* following Definitions 3.4.2 to 3.4.4, but replacing the structure corresponding to \mathbf{Eq} -types with that of \mathbf{Id} -types; as in Theorem 3.4.5, \mathcal{T}_{ITT} is the initial model of ITT. For clarity we rename the concepts defined in Definitions 3.4.2 to 3.4.4 to *model of ETT*, *homomorphism of models of ETT*, and *syntactic model \mathcal{T}_{ETT} of ETT*.

Theorem 4.3.2. *The underlying sets of the syntactic model of ETT support a model of ITT.*

Proof. Intuitively, this means that the syntax of ETT “satisfies the rules of ITT.” Formally, we construct a model \mathcal{M} of ITT whose contexts are the contexts of the syntax of ETT, $\mathbf{Cx}_{\mathcal{M}} := \mathbf{Cx}_{\mathcal{T}_{ETT}}$; whose substitutions are the substitutions of the syntax of ETT, $\mathbf{Sb}_{\mathcal{M}}(\Delta, \Gamma) := \mathbf{Sb}_{\mathcal{T}_{ETT}}(\Delta, \Gamma)$; and likewise for types and terms. For all the rules of ITT that are also present in ETT, we choose the corresponding structure, e.g., $\mathbf{1}_{\mathcal{M}} := \mathbf{1}_{\mathcal{T}_{ETT}}$.

The only subtlety is how to define the \mathbf{Id} -types of \mathcal{M} , and for this we choose the \mathbf{Eq} -types of \mathcal{T}_{ETT} , i.e., $\mathbf{Id}_{\mathcal{M}}(A, a, b) := \mathbf{Eq}_{\mathcal{T}_{ETT}}(A, a, b)$ and $\mathbf{refl}_{\mathcal{M}} := \mathbf{refl}_{\mathcal{T}_{ETT}}$. The reader has already verified in Exercise 4.12 that the \mathbf{J} eliminator is definable in ETT. \square

Corollary 4.3.3. *There is a function $\llbracket - \rrbracket$ that sends contexts (resp., substitutions, types, terms) of ITT to contexts (resp., substitutions, types, terms) of ETT.*

Proof. By Theorem 4.3.2 and the initiality of the syntactic model of ITT, there is a unique homomorphism $f : \mathcal{T}_{ITT} \rightarrow \mathcal{M}$ of models of ITT, and thus in particular there are functions $\mathbf{Cx}_f : \mathbf{Cx}_{\mathcal{T}_{ITT}} \rightarrow \mathbf{Cx}_{\mathcal{M}} = \mathbf{Cx}_{\mathcal{T}_{ETT}}$ and likewise for substitutions, types, and terms. \square

By construction, this translation $\llbracket - \rrbracket$ of ITT to ETT “does nothing” except at **Id**-types, where $\llbracket \mathbf{Id}(A, a, b) \rrbracket = \mathbf{Eq}(\llbracket A \rrbracket, \llbracket a \rrbracket, \llbracket b \rrbracket)$. Intuitively, this is possible because **Eq**-types are defined to have only **refl** as elements, which is strictly stronger than the definition of **Id**-types as “appearing to *other* types to have only **refl** as elements.”

Exercise 4.13. Using Corollary 4.3.3 and Theorem 3.4.8, prove that intensional type theory is consistent.

We can now ask a more precise question:

Exercise 4.14. Suppose that $1 \vdash A$ type in ITT, and that in ETT there is a term $1 \vdash a : \llbracket A \rrbracket$. Then does there necessarily exist a term $1 \vdash a' : A$ in ITT?

By focusing only on types that are well-formed in ITT, this formulation avoids the pitfalls discussed earlier. Perhaps the converse of Exercise 4.14 is more intuitive: *do there exist propositions that can be stated **without** equality reflection, but that can only be proven **with** equality reflection?* Unfortunately, such propositions *do* exist, and thus the answer to Exercise 4.14 is *no*; even worse, the counterexamples are ones that users of type theory are likely to encounter frequently in practice.

Independence We note that counterexamples to Exercise 4.14 are in fact propositions that are *independent* of intensional type theory, i.e., propositions for which neither A nor $A \rightarrow \mathbf{Void}$ are provable—equivalently, *are inhabited*, or contain at least one closed term.

Lemma 4.3.4. *If $1 \vdash A$ type is a counterexample to Exercise 4.14, then A is independent of intensional type theory.*

Proof. By definition, there must exist a term $1 \vdash a : \llbracket A \rrbracket$ in ETT, but no term $1 \vdash a' : A$ in ITT. Thus A is by definition not provable in ITT, so it suffices to show that $A \rightarrow \mathbf{Void}$ is also not provable in ITT. Suppose that there were a term $1 \vdash f : A \rightarrow \mathbf{Void}$ in ITT; then there would also be a term $1 \vdash \llbracket f \rrbracket : \llbracket A \rrbracket \rightarrow \mathbf{Void}$ in ETT, but this would mean there is a closed proof $\llbracket f \rrbracket(a)$ of **Void** in ETT, contradicting its consistency (Theorem 3.4.8). \square

Of course, there are other kinds of independent propositions too; as a sufficiently strong formal system, ITT is subject to Gödel’s incompleteness theorem and thus one can construct independent propositions roughly corresponding to “the type of consistency proofs of ITT.” But for now we restrict our attention to counterexamples to Exercise 4.14, exploring two in particular: function extensionality and uniqueness of identity proofs.

4.3.1 Function extensionality

The principle of *function extensionality* states that for any two functions $f, g : (a : A) \rightarrow B(a)$, if $f(a)$ and $g(a)$ are equal for all $a : A$, then f and g are equal. We reproduce the formal statement of `funext` below, along with its non-dependent special case `funext'`:

$$\text{Funext} = (A : \mathbf{U}) \rightarrow (B : A \rightarrow \mathbf{U}) \rightarrow (f\ g : (a : A) \rightarrow B\ a) \rightarrow \\ ((a : A) \rightarrow \mathbf{Id}(B\ a, f\ a, g\ a)) \rightarrow \mathbf{Id}((a : A) \rightarrow B\ a, f, g)$$

$$\text{Funext}' = (A\ B : \mathbf{U}) \rightarrow (f\ g : A \rightarrow B) \rightarrow \\ ((a : A) \rightarrow \mathbf{Id}(B, f\ a, g\ a)) \rightarrow \mathbf{Id}(A \rightarrow B, f, g)$$

Both of these are counterexamples to Exercise 4.14 and thus independent of ITT. First, we check that $\llbracket \text{Funext} \rrbracket$ is provable in ETT.

Exercise 4.15. Construct a closed term of type $\llbracket \text{Funext} \rrbracket$ in extensional type theory.

Next, we must check that `Funext` is not provable in intensional type theory. As with consistency (Theorem 3.4.7), it suffices to exhibit a model of ITT in which the set of closed terms of type `Funext` is empty. However, it is surprisingly difficult to do so!⁴ One such model is—tautologically—the syntax of ITT itself, or \mathcal{T}_{ITT} ; however, showing that this is the case is precisely what we are already trying to prove. A more useful observation is that the models used to prove normalization contain concrete characterizations of $\text{Tm}(\Gamma, A)$ for all Γ, A and thus it is possible to unfold such a model and explicitly verify that there are no normal forms—and hence no elements whatsoever—of $\text{Tm}(\mathbf{1}, \text{Funext})$ [Hof95a].

Remark 4.3.5. The latter approach is tantamount to the proof-theoretic technique of showing that a formula is not derivable by proving cut elimination for a calculus and then checking by induction that the formula has no cut-free proofs. \diamond

One can also imagine more “mathematical” (and non-initial) models that refute function extensionality. An early example of such a model based on realizability and gluing was given by Streicher [Str93, Chapter 3]; a more recent example is the (categorical) “polynomial” model of von Glehn [Gle14]. In both cases the model construction is somewhat involved but checking that they refute `Funext` is comparatively straightforward. In any case, any of these arguments allows us to conclude:

Theorem 4.3.6. *There is no closed term of type `Funext` in intensional type theory.*

The authors are uncertain to whom this result should be attributed. Turner [Tur89] suggests that it was known to Martin-Löf and it was certainly known to type theorists

⁴There are many simple “countermodels of function extensionality” which fail to validate the η -rule of Π -types and are therefore not models of ITT as we have defined it. They are, however, models of the calculus of inductive constructions, which lacks η for Π -types.

in the 1980s, but the earliest explicit discussion of the independence of Funext we have located is the countermodel of Streicher [Str93].

There are many examples of function extensionality arising in practice. For instance, in ITT we can prove $(n\ m : \text{Nat}) \rightarrow \text{Id}(\text{Nat}, n + m, m + n)$ but not $\text{Id}(\text{Nat} \rightarrow \text{Nat} \rightarrow \text{Nat}, (+), (+) \circ \text{flip})$. Similarly, although $\text{mergeSort}, \text{bubbleSort} : \text{List Nat} \rightarrow \text{List Nat}$ agree on all inputs, we cannot prove they are equal functions. This has real consequences in practice: if we write a function that calls `bubbleSort`, is it equal to the same function where these calls have been replaced by calls to `mergeSort`? If function extensionality held this would follow immediately from `cong`; as it stands, one must manually argue that the text of the function respects swapping subroutines in this way—even though it is impossible to define a function that *doesn't*!

We view the independence of function extensionality as perhaps the greatest failing of intensional type theory, as it frequently causes problems with no benefit,⁵ and it is therefore common to simply *postulate* Funext when working in ITT, that is, to add a rule

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \text{funext} : \text{Funext}} \quad \text{📎}$$

Postulating an axiom in this way is equivalent to prepending every context by a variable of type Funext, and it therefore preserves normalization (a property of *all* contexts) while disrupting canonicity (a property of the *empty* context, which is “no longer empty”).

Exercise 4.16. Argue that postulating Funext causes canonicity to fail. That is, produce a closed term of type **Bool** in ITT adjoined with the above rule that appears to be judgmentally equal to neither **true** nor **false**. (You do not need to formally prove this fact.)

4.3.2 Uniqueness of identity proofs

Our second counterexample to Exercise 4.14 is the principle of *uniqueness of identity proofs* (UIP), which states that any two identifications between the same two terms are themselves identified.

$$\text{UIP} = (A : \mathbf{U}) \rightarrow (a\ b : A) \rightarrow (p\ q : \text{Id}(A, a, b)) \rightarrow \text{Id}(\text{Id}(A, a, b), p, q)$$

In short, UIP asserts that identifications are unique: up to identification, there is at most one proof of $\text{Id}(A, a, b)$ for any $a, b : A$. Types with at most one element are often called (*homotopy*) *propositions*, so we might equivalently phrase UIP as the principle that

⁵There are occasions where one may wish to *not* identify all pointwise-equal procedures, e.g., when studying the runtime of algorithms, but we stress that ITT also does not allow us to *distinguish* pointwise-equal functions; studying runtime in this way requires other axioms and, likely, the removal of β -rules.

propositional equality is a proposition.⁶ Like Funext, UIP is independent of ITT. On the one hand, it holds in ETT and thus cannot be *refuted* by ITT:

Exercise 4.17. Construct a closed term of type $\llbracket \text{UIP} \rrbracket$ in extensional type theory.

To see that UIP is not *provable* in ITT, it again suffices to exhibit a countermodel, a model of ITT in which the set of closed terms of type UIP is empty. The original such countermodel, the *groupoid model of type theory* of Hofmann and Streicher [HS98], is very instructive as well as historically significant as a precursor to homotopy type theory (Chapter 5), so unlike the countermodels of Funext we will take the time to sketch it below.

The groupoid model is similar to the set-theoretic model of type theory (Section 3.5) except that it replaces sets with groupoids, sets equipped with additional structure:

Definition 4.3.7. A *groupoid* $X = (|X|, \mathcal{R}, \text{id}, (-)^{-1}, \circ)$ consists of a set $|X|$, a family of sets \mathcal{R} indexed over $|X| \times |X|$, and dependent functions:

- $\text{id} : \{x : |X|\} \rightarrow \mathcal{R}(x, x)$,
- $(-)^{-1} : \{x y : |X|\} \rightarrow \mathcal{R}(x, y) \rightarrow \mathcal{R}(y, x)$, and
- $(\circ) : \{x y z : |X|\} \rightarrow \mathcal{R}(y, z) \rightarrow \mathcal{R}(x, y) \rightarrow \mathcal{R}(x, z)$,

such that $\text{id} \circ f = f = f \circ \text{id}$, $f \circ f^{-1} = \text{id}$, $\text{id} = f^{-1} \circ f$, and $f \circ (g \circ h) = (f \circ g) \circ h$.

Definition 4.3.8. Given two groupoids X, Y , a *homomorphism of groupoids* $F : X \rightarrow Y$ is a pair of functions $F_0 : |X| \rightarrow |Y|$ and $F_1 : \{x x' : |X|\} \rightarrow \mathcal{R}_X(x, x') \rightarrow \mathcal{R}_Y(F_0(x), F_0(x'))$ for which F_1 commutes with the groupoid operations, i.e.,

- $F_1(\text{id}) = \text{id}$,
- $F_1(f^{-1}) = F_1(f)^{-1}$, and
- $F_1(g \circ f) = F_1(g) \circ F_1(f)$.

Exercise 4.18. For categorically-minded readers: argue that a groupoid is exactly the same as a category all of whose morphisms are isomorphisms, and a homomorphism of groupoids is exactly a functor.

Advanced Remark 4.3.9. The name “groupoid” comes from the perspective that these are a weaker notion of group in which the multiplication is a partial operation. \diamond

⁶The terminology of “propositional equality” is perhaps ill-advised.

We can think of a groupoid as equipping its underlying set with a “proof-relevant notion of equality” which like ordinary equality is reflexive, symmetric, transitive, and respected by functions (groupoid homomorphisms), but unlike ordinary equality “can hold in more than one way.” Following this intuition, we will model closed types $A \in \text{Ty}(1)$ not as sets X but as groupoids $(|X|, \mathcal{R}, \dots)$, closed terms $a \in \text{Tm}(1, A)$ as elements of $|X|$, and closed identifications $p \in \text{Tm}(1, \text{Id}(A, a, b))$ as elements of $\mathcal{R}(a, b)$.

Before outlining the model itself, we give a few examples of groupoids.

Example 4.3.10. Every set A can be regarded as a *discrete groupoid* ΔA in which $\mathcal{R}_{\Delta A}(x, y) = \{\star \mid x = y\}$. The remaining structure is uniquely determined: $\text{id} = \star$, $\star^{-1} = \star$, etc.

Example 4.3.11. Given two groupoids X, Y , the set of groupoid homomorphisms $X \rightarrow Y$ (Definition 4.3.8) admits a natural groupoid structure in which

$$\begin{aligned} \mathcal{R}_{X \rightarrow Y}(F, G) = \\ \{T : (x : |X|) \rightarrow \mathcal{R}(F_0 x, G_0 x) \mid \forall f : \mathcal{R}(x, y). G_1(f) \circ T(x) = T(y) \circ F_1(f)\} \end{aligned}$$

In light of Exercise 4.18, categorically-minded readers might observe that T is exactly a natural transformation from F to G . We leave the remaining structure as an exercise.

Example 4.3.12. For an explicit example of a groupoid that is not discrete, consider the groupoid traditionally called $B(\mathbb{Z}/2)$, whose underlying set is the singleton $\{\star\}$, $\mathcal{R}_{B(\mathbb{Z}/2)}(\star, \star) = \mathbb{Z}/2 = \{0, 1\}$, and the remaining structure is as follows:

$$\begin{aligned} \text{id} &= 0 \\ x \circ y &= x + y \pmod{2} \\ x^{-1} &= x \end{aligned}$$

The reader can check that these operations satisfy the necessary equations. (Hint: this is equivalent to checking that $\mathbb{Z}/2$ with the above id , \circ , and $(-)^{-1}$ forms a group.)

Example 4.3.13. There is a “large” groupoid \mathbb{S} of all “small” sets, where $\mathcal{R}_{\mathbb{S}}(X, Y)$ is the set of bijections between the sets X and Y , and the operations are the identity, inverse, and composition of bijections. This groupoid is not discrete because there can be more than one bijection between a pair of sets, e.g., $\text{id}, \text{swap} \in \mathcal{R}_{\mathbb{S}}(\{\star, \star'\}, \{\star, \star'\})$.

Example 4.3.14. There is a “large” groupoid \mathbb{G} of all “small” groupoids, whose underlying collection is the proper class of all groupoids, and for which $\mathcal{R}_{\mathbb{G}}(X, Y)$ is the set of all groupoid isomorphisms (invertible homomorphisms, or homomorphisms for which F_0 and each F_1 are bijections) from X to Y . The groupoid \mathbb{S} from Example 4.3.13 embeds into \mathbb{G} , so \mathbb{G} is also not discrete.

As in the set-theoretic model of type theory, groupoids and groupoid-indexed families of groupoids form a model of type theory. Writing \mathcal{G} for the groupoid model of (intensional)

type theory and $f : \mathcal{T}_{ITT} \rightarrow \mathcal{G}$ for the homomorphism from the syntactic model to \mathcal{G} , f interprets syntactic contexts Γ as groupoids $Cx_f(\Gamma)$, the closed context $\mathbf{1}$ as the one-element, one-identification groupoid, syntactic substitutions as groupoid homomorphisms, and syntactic types $A \in \text{Ty}(\Gamma)$ as $Cx_f(\Gamma)$ -indexed families of groupoids $(\text{Ty}_f(\Gamma)(A))_{\gamma \in Cx_f(\Gamma)}$. Such a family assigns to each groupoid element $\gamma \in Cx_f(\Gamma)$ a groupoid $(\text{Ty}_f(\Gamma)(A))_\gamma$, and to each identification $\alpha \in \mathcal{R}_{Cx_f(\Gamma)}(\gamma, \gamma')$ a homomorphism $(\text{Ty}_f(\Gamma)(A))_\gamma \rightarrow (\text{Ty}_f(\Gamma)(A))_{\gamma'}$ in a manner compatible with identity and composition. (Using Example 4.3.14, we can repackage the data of such a family quite simply as a groupoid homomorphism $Cx_f(\Gamma) \rightarrow \mathbf{G}$.) Finally, f interprets syntactic terms $a \in \text{Tm}(\Gamma, A)$ as dependent functions assigning to each element $\gamma \in Cx_f(\Gamma)$ of the context an element of the groupoid $(\text{Ty}_f(\Gamma)(A))_\gamma$ in a manner that respects identifications. (We can again phrase this condition as a groupoid homomorphism, but we will not pursue the details further.)

Most of the structure of the groupoid model of type theory mirrors that of the set-theoretic model, with some added complication to account for identifications; for example, rather than interpreting the universe as the large set of all small sets, we interpret it as the large groupoid \mathbf{G} of all small groupoids (Example 4.3.14). The key departure is in the interpretation of **Id**-types: for closed $A \in \text{Ty}_{\mathcal{G}}(\mathbf{1}_{\mathcal{G}})$ and $a, b \in \text{Tm}_{\mathcal{G}}(\mathbf{1}_{\mathcal{G}}, A)$, the \mathcal{G} -identity type $\mathbf{Id}_{\mathcal{G}}(A, a, b)$ is chosen to be (the discrete groupoid on) the set of identifications in the groupoid A between a and b , namely $\Delta\mathcal{R}_A(a, b)$.

It is not at all obvious that such an interpretation supports **J**, but this is the force of the groupoid model: because all types and terms respect identifications, it is in fact the case that dependent functions from **Id**-types into any \mathcal{G} -type are generated by the data of where to send **refl**. Interested readers can find these and all the other details in the paper of Hofmann and Streicher [HS98].

Theorem 4.3.15 (Hofmann and Streicher [HS98]). *There is no closed term of type **UIP** in intensional type theory.*

Proof. This follows immediately from the fact that the groupoid model interprets **UIP** as the empty groupoid, whose proof we sketch below. Recall that:

$$\mathbf{UIP} = (A : \mathbf{U}) \rightarrow (a \ b : A) \rightarrow (p \ q : \mathbf{Id}(A, a, b)) \rightarrow \mathbf{Id}(\mathbf{Id}(A, a, b), p, q)$$

A term of this type in \mathcal{G} would be a dependent function out of the interpretation of \mathbf{U} , which is the groupoid of groupoids \mathbf{G} . Suppose that such a function exists; then we could apply it to the groupoid $B(\mathbb{Z}/2) \in \mathbf{G}$ defined in Example 4.3.12, then twice to the unique element $\star \in |B(\mathbb{Z}/2)|$ of that groupoid, and then to the two distinct identifications $0, 1 \in \mathcal{R}_{B(\mathbb{Z}/2)}(\star, \star)$. The result would have to be a proof that $0 = 1$, which is false. \square

4.3.2.1 Towards homotopy type theory

The busy reader may wish to skip this section initially, although we suggest that they return to it when reading Section 5.3. The groupoid model demonstrates that **Id**-types support richer interpretations than merely equations: identifications can be any data that is respected by all the constructs of type theory.

Although the groupoid model provides us with interesting examples of identity types, we note that the identity types of any groupoid X , $\Delta\mathcal{R}_X(x, y)$, are always discrete groupoids with no interesting identifications of their own. Thus the groupoid model *does* validate the “uniqueness of identity proofs of identity proofs”:

$$\begin{aligned} \text{UIPIP} = & (A : \mathbf{U}) \rightarrow (a\ b : A) \rightarrow (p\ q : \mathbf{Id}(A, a, b)) \rightarrow \\ & (\alpha\ \beta : \mathbf{Id}(\mathbf{Id}(A, a, b), p, q)) \rightarrow \mathbf{Id}(\mathbf{Id}(\mathbf{Id}(A, a, b), p, q), \alpha, \beta) \end{aligned}$$

Like UIP, this principle is also independent of ITT, and we can construct a countermodel in *2-groupoids*, which contain a second level of “2-identifications” $\mathcal{R}^2(p, q)$ between any pair of identifications $p, q \in \mathcal{R}(a, b)$ between elements a, b . Although we will not define these precisely, we note that the passage from groupoids to 2-groupoids is analogous to the passage from sets to groupoids; for instance, every groupoid can be regarded as a discrete 2-groupoid with the same elements and 1-identifications but with trivial 2-identifications.

The story once again repeats for the 2-groupoid model of type theory, and in fact for any n : there is a model of ITT in which closed types are interpreted as n -groupoids, and this model refutes $\text{U}(\text{IP})^n$ but validates $\text{U}(\text{IP})^{n+1}$. In fact, this suggests correctly that ordinary groupoids ought to be thought of as *1-groupoids* and sets as *0-groupoids*; indeed, the set (0-groupoid) model of type theory validates UIP^1 . Looking downward, the large 0-groupoid of (-1) -groupoids is the *set of propositions* $\{\emptyset, \{\star\}\}$ as we will see in Section 5.3.1.

But what about for *all* n ? Is it possible to construct a model that simultaneously refutes $\text{U}(\text{IP})^n$ for every $n \in \mathbb{N}$? Intuitively, such a model would have to interpret closed types as “ ∞ -groupoids” with countably infinite towers of identifications. The answer is *yes* [War08, Corollary 4.26], and in fact Voevodsky’s simplicial model of homotopy type theory [KL21] can be seen as precisely such a model [KS15].

4.3.3 Hofmann’s conservativity theorem

We have generated an infinite stream of counterexamples to Exercise 4.14—propositions that are provable in ETT but not ITT—namely Funext and $\text{U}(\text{IP})^n$ for $n \geq 1$. Is there a third class of counterexamples? Surprisingly, *no*: all counterexamples to Exercise 4.14 are generated by Funext and UIP in a precise sense. (Note that UIP implies $\text{U}(\text{IP})^n$ for $n > 1$.)

To state this claim more precisely, let us write

$$\Gamma_{ax} := \mathbf{1}, \text{funext} : \text{Funext}, \text{uip} : \text{UIP}$$

for the ITT context containing two variables, one of type Funext and one of type UIP ; types and terms of ITT in context Γ_{ax} are in bijection with closed types and terms of intensional type theory extended by two rules postulating Funext and UIP . Then, Hofmann’s celebrated *conservativity result* states that:

Theorem 4.3.16 (Hofmann [Hof95a]). *Suppose that $\Gamma_{ax} \vdash A$ type in ITT, and $\llbracket \Gamma_{ax} \rrbracket \vdash a : \llbracket A \rrbracket$ in ETT; then there exists a term $\Gamma_{ax} \vdash a' : A$ in ITT.*

In Exercises 4.15 and 4.17 the reader has constructed proofs $\mathbf{1} \vdash p : \llbracket \text{Funext} \rrbracket$ and $\mathbf{1} \vdash q : \llbracket \text{UIP} \rrbracket$ of function extensionality and UIP in ETT, so we can discharge the hypotheses of $\llbracket \Gamma_{ax} \rrbracket$ to obtain the following corollary:

Corollary 4.3.17. *If $\Gamma_{ax} \vdash A$ type in ITT and $\mathbf{1} \vdash a : \llbracket A \rrbracket [p/\text{funext}, q/\text{uip}]$ in ETT, then there exists a term $\Gamma_{ax} \vdash a' : A$ in ITT.*

Corollary 4.3.17 is great news: although ITT is weaker than ETT, it is weaker by exactly two principles, namely function extensionality and uniqueness of identity proofs. We are led naturally to wonder whether there is a “best of both worlds”:

Exercise 4.19. Can we extend intensional type theory (with new terms and/or equations) in such a way that Funext and UIP are derivable, and the resulting type theory enjoys both canonicity *and* normalization?

If we are satisfied with only *one of* canonicity or normalization, note that ETT is such an extension of ITT satisfying canonicity (Theorem 3.4.12) but not normalization (Section 3.6); on the other hand, extending ITT with axioms for Funext and UIP trivially makes these provable and satisfies normalization (Theorem 4.2.4) but not canonicity (Exercise 4.16).

Remark 4.3.18. Such tradeoffs are common in the design of type theory: canonicity says that a type theory has “enough” equations, whereas normalization generally requires that there not be “too many”; it can be hard to find the right balance. \diamond

Type theorists have considered Exercise 4.19 since the 1990s, and there is some good news to report. If we are content for the moment to solve only the problem of UIP (ignoring Funext), there is in fact a rather modest extension of ITT that satisfies canonicity and normalization and in which UIP is provable.

For this, it will help us to consider an equivalent formulation of UIP due to Streicher [Str93] known as *Axiom K*:⁷

$$K = (A : \mathbf{U}) \rightarrow (a : A) \rightarrow (p : \text{Id}(A, a, a)) \rightarrow \text{Id}(\text{Id}(A, a, a), p, \text{refl})$$

⁷In light of Remark 4.2.3, perhaps the reader can guess where the name K comes from.

It is easy to see that \mathbf{K} follows from \mathbf{UIP} , as it is the special case of \mathbf{UIP} in which a and b are the same and one of the identity proofs is \mathbf{refl} . The other direction of the bimplication is more subtle, but follows from a careful application of \mathbf{J} , or identity elimination.

Exercise 4.20. Prove that \mathbf{K} implies \mathbf{UIP} in \mathbf{ITT} .

As with \mathbf{subst} and \mathbf{uniq} , there is a sensible definitional equality with which to equip $\mathbf{k} : \mathbf{K}$, namely $\mathbf{k} A a \mathbf{refl} = \mathbf{refl}$, and we can even rephrase \mathbf{k} as a “second elimination principle” of \mathbf{Id} -types as follows:

$$\frac{\Gamma \vdash a : A \quad \Gamma \vdash p : \mathbf{Id}(A, a, a) \quad \Gamma.A.\mathbf{Id}(A[\mathbf{p}], \mathbf{q}, \mathbf{q}) \vdash B \text{ type} \quad \Gamma.A \vdash b : B[\mathbf{id.refl}]}{\Gamma \vdash \mathbf{K}(b, p) : B[\mathbf{id.a.p}]}$$

$$\frac{\Gamma \vdash a : A \quad \Gamma.A.\mathbf{Id}(A[\mathbf{p}], \mathbf{q}, \mathbf{q}) \vdash B \text{ type} \quad \Gamma.A \vdash b : B[\mathbf{id.refl}]}{\Gamma \vdash \mathbf{K}(b, \mathbf{refl}) = b[\mathbf{id.a.refl}] : B[\mathbf{id.a.refl}]}$$

$$\frac{\Gamma \vdash a : A \quad \Gamma \vdash p : \mathbf{Id}(A, a, a) \quad \Delta \vdash \gamma : \Gamma \quad \Gamma.A.\mathbf{Id}(A[\mathbf{p}], \mathbf{q}, \mathbf{q}) \vdash B \text{ type} \quad \Gamma.A \vdash b : B[\mathbf{id.refl}]}{\Delta \vdash \mathbf{K}(b, p)[\gamma] = \mathbf{K}(b[(\gamma \circ \mathbf{p}).\mathbf{q}], p[\gamma]) : B[\gamma.a[\gamma].p[\gamma]}}$$

It is instructive to compare the rules for \mathbf{K} to those of \mathbf{J} , whose motives

$$\Gamma.A.A[\mathbf{p}].\mathbf{Id}(A[\mathbf{p}^2], \mathbf{q}[\mathbf{p}], \mathbf{q}) \vdash C \text{ type}$$

quantify over *both* sides of the identification. Although \mathbf{J} may seem superficially more general, neither \mathbf{J} nor \mathbf{K} imply the other. On the one hand, \mathbf{K} is equivalent to \mathbf{UIP} , which is independent of \mathbf{ITT} ; on the other hand, we needed the additional flexibility of \mathbf{J} to define \mathbf{subst} (Lemma 4.2.6), and we invite the reader to attempt this definition with \mathbf{K} alone.

Although adding the above rules for \mathbf{K} to intensional type theory breaks the pattern of inductive types we established in Section 2.5, the resulting theory continues to enjoy all the good properties of intensional type theory.

Theorem 4.3.19. *Intensional type theory plus the above rules for \mathbf{K} satisfies consistency, canonicity, normalization, has invertible type constructors, and also validates \mathbf{UIP} .*

In fact, \mathbf{K} was originally introduced not to restore extensionality to \mathbf{ITT} but in the study of *dependent pattern-matching*, where early formulations of pattern-matching for dependent type theory [Coq92] were found to derive \mathbf{K} and were thus stronger than the standard rules of \mathbf{ITT} . Although researchers have subsequently formulated a weaker notion of pattern-matching that does not derive \mathbf{K} [CDP14], many proof assistants such as Agda still include \mathbf{K} by default, often via pattern-matching.

Unfortunately it is significantly more challenging to add function extensionality to ITT in a satisfactory (canonicity-preserving) fashion, either in tandem with or independently of \mathbf{K}/UIP . There are a number of type theories that admit function extensionality and satisfy all the relevant metatheorems, most notably *observational type theories* (Section 4.4, which also validate UIP) and *cubical type theories* (Chapter 5, which intentionally do not validate UIP), but these systems are quite a bit more complex than ITT and have not supplanted it.

Thus, despite its shortcomings, many practitioners choose to work in ITT extended with an axiom for function extensionality and either an axiom for UIP or a version of dependent pattern-matching that validates \mathbf{K} .

4.4[★] *Observational type theory* (DRAFT)

Further reading

We have mentioned previously that proof assistants decide equality of terms using a type-sensitive algorithm known as normalization by evaluation (NbE). Proofs of the normalization metatheorem for intensional type theory proceed by establishing that NbE is sound and complete for the equational theory of ITT, using a proof technique known as Kripke logical relations. There are many papers dedicated to proving normalization for variants of ITT; Abel [Abe13] includes a lengthy exposition starting with the non-dependent case, Abel, Öhman, and Vezzosi [AÖV17] formalize their proof in Agda, and Coquand [Coq19] and Sterling [Ste21] present semantic formulations of NbE that are significantly less technical but require more mathematical sophistication.

As for the independence and conservativity theorems discussed in Section 4.3, the theses of Streicher [Str93] and Hofmann [Hof95a] remain excellent references; however, a more modern account is available in the thesis of Winterhalter [Win20], and recent advances in semantics have enabled much shorter albeit sophisticated proofs of conservativity [KL23].

The independence of function extensionality from ITT has led to a cottage industry of observational type theories as discussed in Section 4.4; the authors are biased but recommend Sterling, Angiuli, and Gratzer [SAG22, Section 1] for a brief history of equality in type theory. On the other hand, the independence of UIP has spawned an entire *subdiscipline*, homotopy type theory (Chapter 5). Models of homotopy type theory, such as Voevodsky’s simplicial model [KL21], can be seen as vast generalizations of the groupoid model of Hofmann and Streicher [HS98].

Univalent type theories (DRAFT)

In Chapter 4 we studied the intensional identity type better-behaved replacement for Eq which satisfies canonicity, normalization, and decidable type-checking. Unfortunately, these extra metatheorems came a cost: sensible results such as UIP and function extensionality were independent of ITT. This led to the question of whether it was possible to extend ITT to a theory satisfying canonicity, normalization, and decidable type-checking while also validating both UIP and function extensionality (Section 4.4).

In this chapter we survey the a more radical proposal. Rather than constructing a type theory extending ITT to better capture ETT, we consider extending type theory with a reasoning principle *incompatible* with UIP: Voevodsky’s univalence axiom. Intensional type theory extended with univalence, homotopy type theory (HoTT) is therefore incompatible with ETT but features new reasoning principles to compensate for this loss.

We shall explore some of the remarkable consequences of univalence and the new reasoning principles it adds to type theory. More than this, we will how that univalence is an inevitable consequence of attempting to take seriously our insistence that types be characterized uniquely by their behavior. In particular, univalence follows from attempting to equip the universe with a good “mapping-in” property akin to that enjoyed by Π and Σ .

On first inspection, HoTT brings us no closer to our overarching goal of finding a type theory which supports both powerful reasoning principles along with a large suite of metatheorems. Indeed, HoTT is roughly defined by adding the univalence axiom to ITT, the type theory satisfies normalization but fails to validate canonicity.

We resolve this issue with the introduction of a more refined type theory with univalence: cubical type theory. We shall see how cubical type theory incorporates univalence and its consequences while managing to satisfy canonicity, normalization, and decidable type-checking. This balance is achieved by reimagining the intensional identity type such that it can be characterized by a “mapping-in” property just as was done with Eq .

These topics bring the reader all the way present day research in type theory. Univalence was first introduced in 2010 [Voe10], much of the material on homotopy type theory we discuss was put forward in 2013 [UF13], and cubical type theory dates to 2018 [CCHM18] with key results only proven in 2021 [SA21]. Moreover, the story of univalent type theories is emphatically ongoing. Researchers presently attempting to produce other versions of type theory validating univalence as well as canonicity and normalization [Shu23]. Even the consequences simply axiomatizing univalence within ITT is the subject of active study.

In this chapter In Sections 5.1 and 5.2, we motivate, define, and explore univalence as an axiom within intensional type theory. We begin with the simpler form of propositional univalence in Section 5.1 and, with this intuition to hand, introduce full univalence in Section 5.2. There we survey some of the highlights of homotopy type theory and show, in particular, that its inclusion invalidates UIP. In Section 5.3 show how univalence arises naturally from attempting to give universes a mapping-in property and sketch the semantic model proving the consistency of univalence. Finally, in Section 5.4 we discuss the key ideas behind cubical type theory and explain the details and implications in Section 5.5.

Goals of the chapter By the end of this chapter, you will be able to:

- Define the univalence axiom as well as its key constituents such as `isEquiv`.
- Work with core concepts of HoTT such as `h-levels` and higher inductive types.
- Sketch how univalence relates to a mapping-in property for the universe
- Describe how (∞) -groupoids arise in models of univalence.
-

5.1 Propositions in intensional type theory

The central topic of this chapter is univalence, which states that the identity type when applied to elements of the universe $\text{Id}(U, A, B)$ is equivalent to equivalences between A and B . However, univalence has a major deficiency: unlike other principles and axioms we have encountered, it is *not* compatible with viewing types as mere sets.

Before moving on to full univalence, we therefore discuss a halfway point: propositional univalence. This axiom is a consequence of full univalence obtained by restricting to types A, B which have at most one element. Crucially, propositional univalence *is* compatible with set-theoretic intuitions and thus we can give a more precise account of how it is justified without pitching the reader into any homotopy theory. Moreover, propositional univalence is a useful reasoning principle even in the context of extensional type theory and thus of independent interest.

For this section, we shall assume that the reader has some familiarity with the notion of a *proposition* in type theory. We refer them to Section 2.8 for further details.

Notation 5.1.1. In this section, we return to informal intensional type theory as used in Chapter 1 and Section 4.1.

Assumption 5.1.2. We will also assume that function extensionality holds in this section and, in particular, that there is a constant `funext` : `Funext` (see Section 4.3 for the notation).

5.1.1 Homotopy propositions

What, exactly, is the difference between an arbitrary type and a type which ought to be regarded as a proposition? As discussed in Section 2.8: a proper type may have distinct inhabitants, but propositions have at most one element. In extensional type theory, we crystallized this by requiring that if $a, b : A$ then $a = b : A$ definitionally. Equivalently, we could ask that if $a, b : A$ then $\text{Eq}(A, a, b)$ is inhabited. In intensional type theory these two approaches no longer coincide: we can either require that a proposition A is *definitionally irrelevant*—any two inhabitants are definitionally equal—or merely ask that they be *propositionally irrelevant* such that for any two elements a, b there is an element $\text{Id}(A, a, b)$. Customarily, one refers to the first variety of propositions as *strict propositions* and the second as *homotopy propositions*.

Without equality reflection, strict and homotopy propositions do not coincide. Worse, neither is a clearly superior choice and both have advantages and disadvantages. Strict propositions are more ergonomic and do not require the user to provide proofs equating two elements which must, by construction, be equal [Gil+19]. On the other hand, homotopy propositions arise more naturally in type theories without equality reflection. For instance, given a list of natural numbers $l : \text{List Nat}$ the type of minimal elements is a homotopy proposition:

$$\text{MinimumElements}(l) = \sum_{n:\text{Nat}} n \in l \times (m : \text{Nat}) \rightarrow m \in l \rightarrow n \leq m$$

This type, however, is not a strict proposition in ITT (even if l is empty!) owing to the absence of η laws for Nat and Void . More generally, we can consider $A : \mathbf{U}$ and use the full power of type theory to *prove* that A is a homotopy proposition. It is only in the most trivial of cases where such an A will be a strict proposition. Within this chapter we shall encounter various types which are homotopy propositions for subtle reasons and this flexibility is then crucial. For this reason, among others, we shall focus exclusively on homotopy propositions in this chapter.

We may define a predicate $\text{isHProp} : \mathbf{U} \rightarrow \mathbf{U}$ capturing whether a given type is a homotopy proposition (see Exercise 2.38). Consequently, we can define the universe of homotopy propositions without extending our theory at all:

$$\text{HProp}_i := \sum_{A:\mathbf{U}_i} \text{isHProp } A$$

Remark 5.1.3. In truth, there is a final and most important reason for us to use homotopy propositions: they scale properly to homotopy type theory while strict propositions simply do not. However, this claim is by nature difficult to substantiate before we introduce homotopy type theory. \diamond

Notation 5.1.4. We shall usually suppress the index on HProp_i and simply write HProp . Furthermore, we generally treat the projection $\text{HProp} \rightarrow \mathbf{U}$ as silent. We shall also write $A : \text{HProp}$ if $A : \mathbf{U}$ to signify that there exists some term ϕ such that $(A, \phi) : \text{HProp}$.

Exercise 5.1. Suppose that $B : A \rightarrow \mathbf{HProp}$, show that $(a : A) \rightarrow B a : \mathbf{HProp}$. (Hint: use `funext`.)

Our main goal in this section is not, however, to extol the virtues of \mathbf{HProp} . In fact, we want to point out a key deficiency with this type if it is really to represent “a universe of propositions”: it has far, far too many inhabitants according to \mathbf{Id} . In particular, a pair of `hprops` A, B will hardly ever be identifiable even if they represent the same proposition. Consider, for instance, `Unit` and `Unit × Unit`. These are not equal within \mathbf{U} and so cannot be equal in \mathbf{HProp} . What, however, is to be gained by distinguishing them? Essentially nothing.

The situation is similar to function extensionality in Chapter 4; we are not allowed to identify these two types, but nothing within our theory can treat them differently. Indeed, really our only means of interacting with elements of \mathbf{HProp} (or, indeed, elements of \mathbf{U}) is to attempt to construct elements of these types and to ask whether these elements may be identified. For this purpose, `Unit` and `Unit × Unit` are completely interchangeable: both are inhabited and have exactly one element.

As it stands, \mathbf{Id} measures roughly whether two propositions have been built using an identical sequence of type formers. However, no meaningful operation on propositions depends on this fact! Indeed, if we informally regard propositions in the set-theoretic model, we find that we are splitting hairs over whether e.g., $\{\star\}$ and $\{(\star, \star)\}$ are equal. They are certainly different sets, but they also clearly encode the same proposition (\top) and it would therefore be pleasant if our notion of equality captured this fact.

To further belabor the point, let us import a reasoning principle from the set model of type theory: the law of the excluded middle (LEM). This principle states that given $\phi : \mathbf{HProp}$ either ϕ or $\neg\phi$ is inhabited. This is certainly not derivable in type theory and we do not suggest its addition, but we do note that if it holds then we can define a map:

$$\text{isTrue} : \mathbf{HProp} \rightarrow \mathbf{Bool}$$

This map actually gives rise to something like a bijection between \mathbf{HProp} and \mathbf{Bool} with the inverse being $\mathbf{Id}(\mathbf{Bool}, \text{true}, -)$. It is only “like” a bijection, because while $\text{isTrueId}(\mathbf{Bool}, \text{true}, \text{true})$ holds, $\mathbf{Id}(\mathbf{Bool}, \text{isTrue}(\phi), \text{true})$ is merely logically equivalent to ϕ rather than equal.

The idea behind our new reasoning principle—*propositional univalence*—is to rectify this deficiency by postulating that logically equivalent propositions are identifiable.

$$\begin{aligned} (\leftrightarrow) : \mathbf{U} &\rightarrow \mathbf{U} \rightarrow \mathbf{U} \\ A \leftrightarrow B &= (A \rightarrow B) \times (B \rightarrow A) \end{aligned}$$

$$\text{propUnivalence} : (A B : \mathbf{HProp}) \rightarrow (A \leftrightarrow B) \rightarrow \mathbf{Id}(\mathbf{HProp}, A, B)$$

In fact, we can derive a seemingly stronger result from `propUnivalence`: the identity type $\mathbf{Id}(\mathbf{HProp}, A, B)$ is actually *isomorphic* to $A \leftrightarrow B$.

Notation 5.1.5. We write $A \cong B$ for the type following type:

$$(f : A \rightarrow B) \times (g : B \rightarrow A) \times \mathbf{Id}(A \rightarrow A, g \circ f, \mathbf{id}) \times \mathbf{Id}(B \rightarrow B, f \circ g, \mathbf{id})$$

Theorem 5.1.6. `propUnivalence` induces an isomorphism $(A \leftrightarrow B) \cong \mathbf{Id}(\mathbf{HProp}, A, B)$

Proof. We can easily define a putative inverse `propUnivalence`⁻¹ to `propUnivalence`: given an identification $p : \mathbf{Id}(\mathbf{HProp}, A, B)$ we define a biimplication $A \leftrightarrow B$ using the tuple $(\text{subst } \mathbf{id} \ p, \text{subst } \mathbf{id} \ (\text{sym } p))$. Moreover, using Exercise 5.2 we know that `propUnivalence`⁻¹ \circ `propUnivalence` can be identified with `id`.

For the reverse, we note that to show that $i = \text{propUnivalence} \circ \text{propUnivalence}^{-1}$ is an idempotent map: there is an identification between $i \circ i$ and i . However, any idempotent map $\mathbf{Id}(A, a, b) \rightarrow \mathbf{Id}(A, a, b)$ must be the identity (Exercise 5.3) completing the proof. \square

This result leads us to a reformulation of propositional univalence:

Corollary 5.1.7. *The canonical map $\mathbf{Id}(\mathbf{HProp}, A, B) \rightarrow (A \leftrightarrow B)$ induced by `subst` is an isomorphism.*

Exercise 5.2. If $A, B : \mathbf{HProp}$, then $A \leftrightarrow B : \mathbf{HProp}$.

Exercise 5.3. Prove the following result due to Escardó [Esc14]: if $i : \mathbf{Id}(A, a, b) \rightarrow \mathbf{Id}(A, a, b)$ is idempotent, it is equal to `id`. (Hint: argue first using the elimination principle that $i(p)$ can be identified with `trans (i refl) p` and then use this to identify `trans (i refl) (i refl)` with `(i refl)`).

Returning to our earlier discussion regarding `HProp` and `Bool` in the presence of LEM, we might alternatively define `propUnivalence` in this setting by stating that the canonical map `Bool` \rightarrow `HProp` is a genuine isomorphism. This insight leads us to a construction of a model of type theory in sets which validates propositional univalence:

Theorem 5.1.8. *Both intensional and extensional type theory with propositional univalence is consistent.*

We shall return to the semantics of propositional univalence more thoroughly in Section 5.3. One takeaway from the above consistency result, however, is that `propUnivalence` is compatible with our existing intuitions for type theory: if types are sets, there is no particular harm in identifying propositions up to logical equivalence.

Towards univalence All the motivations that led us to desiring propositional univalence, however, are not limited to mere propositions. We only interact with types by studying their inhabitants and so we may just as well ask that two elements of \mathbf{U} be identified when they are isomorphic. Hence, we are naturally led to wonder about what happens if we drop the requirement that propUnivalence applies only to propositions. That is, what the status of the full univalence axiom:

$$(A B : \mathbf{U}) \rightarrow \text{isEquiv}(\mathbf{Id}(\mathbf{U}, A, B) \rightarrow A \simeq B)$$

Here we have deliberately introduced some undefined notation: isEquiv and \simeq rather than our previous \cong and “is an isomorphism”. We shall see presently that these replacements are necessary and that a more refined notion of invertability is required for full univalence than for its propositional cousin. In fact, this is only the beginning of the subtlety. We shall see that, unlike propositional univalence, full univalence is incompatible with many of our set-theoretic intuitions for \mathbf{Id} and that univalence *contradicts* UIP as well as its weaker cousins $\mathbf{U}(\text{IP})^n$. Thus, while both are natural, propositional univalence is a relatively anodyne extension compatible with standard interpretations of type theory. On the other hand, univalence leads us to a far richer theory with an entirely different character than extensional type theory.

5.2 Homotopy type theory

In this section, we arrive at the core topic of this chapter: univalence. Having discussed the simpler case of propositional univalence, we introduce the full univalence axiom and, with it, core homotopy type theory: the extension of ITT with the univalence axiom. In order to properly define this, we must give the precise type of the univalence axiom; a more subtle endeavor than one might guess! With univalence to hand, we investigate its most basic consequences and notions such as the homotopy levels of types. We also introduce one of the most interesting applications of homotopical thinking: higher inductive types. We show how HoTT suggests a novel form of inductive type which allows us to describe quotients, propositional truncations, and even topological-inspired objects in a straightforward conceptual manner.

Notation 5.2.1. In this section, we return to informal intensional type theory as used in Chapter 1 and Section 4.1.

Assumption 5.2.2. We will also assume that function extensionality holds in this section and, in particular, that there is a constant $\text{funext} : \text{Funext}$ (see Section 4.3 for the notation).

Related reading Thus far we have limited ourselves to describing additional reading in supplementary sections at the end of this chapter. This was largely because most other sources did not quite match the goals we set ourselves in this book. Fortunately, there are multiple book-length treatments of homotopy type theory freely available and we encourage the reader to consult them. In particular, the *HoTT Book* [UF13] and Rijke’s forthcoming book [Rij22] are both excellent introductions to the subject. There are also multiple formalizations of the core contents of homotopy type theory in proof assistants [VAG+20; Esc+14; Rij+21] and these may also be useful to consult.

We will endeavor to provide citations for the related results in Univalent Foundations Program [UF13] to expedite referencing the HoTT book. The reader will also find Chapters 9 through 17 of Rijke [Rij22] (where chapters are relatively short) to contain a superset of the material discussed in this section.

Remark 5.2.3. In general, the reader who has made it to this point now has the necessary prerequisites to read a great deal of the type theory literature. We encourage them to do so; the goal of this book is to prepare a student to engage with the literature and this is best done through practice! \diamond

5.2.1 The Univalence Axiom

We have begun this section with the rough outline of the univalence axiom in mind: the transport function $\text{transp} = \text{subst id} : \mathbf{Id}(U, A, B) \rightarrow (A \rightarrow B)$ is an injection whose image is precisely the subtype of equivalences from A to B , $A \simeq B$. Our goal is now to nail down the relevant definitions here. We must first explain what it means for a map to be an equivalence *i.e.*, define a proposition $\text{isEquiv}(f)$ (Theorem 5.2.5). If we then define the subtype of equivalences $A \simeq B = \sum_{f:A \rightarrow B} \text{isEquiv}(f)$, we must then show that transp can be written as the composite $\text{fst} \circ \text{idtoequiv}$ where idtoequiv is the (necessarily unique) map $\mathbf{Id}(U, A, B) \rightarrow (A \simeq B)$ (Lemma 5.2.8). Finally, we can then formally state univalence by requiring that idtoequiv is itself an equivalence (Definition 5.2.9).

Without further ado, let us consider the first point: how should we define $\text{isEquiv}(f)$? A first definition might follow what we would say in set theory: $f : A \rightarrow B$ is an equivalence if it has an inverse:

$$\begin{aligned} \text{hasTwoSidedInverse}(f) = \\ \sum_{g:B \rightarrow A} (\mathbf{Id}(A \rightarrow A, g \circ f, \text{id})) \times (\mathbf{Id}(B \rightarrow B, f \circ g, \text{id})) \end{aligned}$$

In this case, $A \simeq B$ would simply become $A \cong B$ as specified above.

Exercise 5.4. Construct an element of $\text{idIsIso} : \text{hasTwoSidedInverse}(\text{id})$.

This definition certainly correlates with our expectations but, surprisingly, it is simply incorrect and asserting univalence based on this definition will result in an inconsistent system. The root of the issue is following metatheorem:

Theorem 5.2.4. *The proposition $\text{isHProp}(\text{hasTwoSidedInverse}(f))$ is independent of ITT.*

In fact, with univalence we will be able to *refute* $\text{isHProp}(\text{hasTwoSidedInverse}(f))$. In such a situation the type $A \simeq B = \sum_{f:A \rightarrow B} \text{hasTwoSidedInverse}(f)$ is not a subtype of $A \rightarrow B$. While this defies our intuitions for $A \simeq B$, it is actually more serious than that. Essentially, asserting univalence allows us to convert $p : A \simeq B$ into an identification $\text{Id}(U, A, B)$ and then—using **J**—pretend that that identification is reflexivity. All told, if we are essentially able to assume that (1) $A = B$ (2) $p = (\text{id}, \text{idIsIso})$ when constructing a term depending on p . As noted above, however, univalence also gives us the ability to construct an element of $\text{hasTwoSidedInverse}(\text{id})$ which is *provably distinct* from idIsIso and this leads to a contradiction.

The upshot of this is that we must find a more refined definition of the proposition isEquiv which is, in particular, a proposition. We emphasize that we still want it to be the case that if $\text{isEquiv}(f)$ holds then f has an inverse. We just cannot have $\text{isEquiv}(f)$ by the literal data of such a two-sided inverse because the actual choice of a particular inverse encodes too much data and leads to inconsistencies. In summary, our specification for $\text{isEquiv}(f)$ is the following three points:

1. A map $\text{hasTwoSidedInverse}(f) \rightarrow \text{isEquiv}(f)$.
2. Another map $\text{isEquiv}(f) \rightarrow \text{hasTwoSidedInverse}(f)$.
3. $\text{isEquiv}(f) : \text{HProp}$

It turns out that any definition of isEquiv satisfying these three properties is as good as any other. The following exercise makes this claim precise:

Exercise 5.5. Show that if $\text{isEquiv}_1(f)$ and $\text{isEquiv}_2(f)$ both satisfy the above constraints, then there is a map $\alpha : \text{isEquiv}_1(f) \rightarrow \text{isEquiv}_2(f)$ such that $\text{hasTwoSidedInverse}(\alpha)$ holds.

Arguably, we could stop our discussion of isEquiv here and simply insist that some definition of isEquiv exists. However, for the sake of completeness we present a valid definition of isEquiv and refer the reader to Univalent Foundations Program [UF13, Chapter 4] for an exhaustive discussion.

Theorem 5.2.5. *If $f : A \rightarrow B$ then the following is a valid definition of $\text{isEquiv}(f)$:*

$$\begin{aligned} \text{fib}(f, b) &= \sum_{a:A} \mathbf{Id}(B, f a, b) \\ \text{isContr}(X) &= \sum_{x:X} (y : X) \rightarrow \mathbf{Id}(X, x, y) \\ \text{isEquiv}(f) &= (b : B) \rightarrow \text{isContr}(\text{fib}(f, b)) \end{aligned}$$

Remark 5.2.6. In the above, $\text{fib}(f, b)$ is the preimage $f^{-1}(b)$ or *fiber* of f over b while $\text{isContr}(X)$ states that X is a type with precisely one element. Informally then, $\text{isEquiv}(f)$ just when each preimage $f^{-1}(b)$ has exactly one element. In homotopical parlance, f is an equivalence just when it has contractible fibers. \diamond

To prove this theorem, we require the following fact from the HoTT book.

Lemma 5.2.7 (Lemma 3.11.4 [UF13]). *For any type A , $\text{isContr}(A) : \mathbf{HProp}$.*

Proof of Theorem 5.2.5. We have three things to prove.

First, we must show that if $f : A \rightarrow B$ has a two-sided inverse $g : B \rightarrow A$, then f has contractible fibers. Let us then fix $b : B$. We must show the following to be contractible:

$$\sum_{a:A} \mathbf{Id}(B, f a, b)$$

Using the assumption that g is a two-sided inverse, we note that it suffices to show that $\sum_{a:A} \mathbf{Id}(A, g(f a), g b)$ is contractible. To show this, it suffices $\sum_{a:A} \mathbf{Id}(A, a, g b)$ but this last type is easy to inhabit as we may choose $(g b, \mathbf{refl})$. We must then show exhibit an identification between $(g b, \mathbf{refl})$ and an arbitrary element $p : \sum_{a:A} \mathbf{Id}(A, g(f a), g b)$. While we omit the details, this identification is almost exactly defined by $\mathbf{snd}(p)$.

Second, we must show that if $p : \text{isEquiv}(f)$ implies $\text{hasTwoSidedInverse}(f)$. Let us define a putative inverse g to f as follows:

$$g b = \text{fst}(\text{fst}(p b))$$

Explicitly, g is defined to be the composite of the following chain of operations:

$$\begin{aligned} b : B & \\ \mapsto p b : \text{isContr}(\text{fib}(f, b)) & \\ \mapsto \text{fst}(p b) : \text{fib}(f, b) = \sum_{a:A} \mathbf{Id}(B, f a, b) & \\ \mapsto \text{fst}(\text{fst}(p b)) : A & \end{aligned}$$

It is then easy to define a function $\beta : (b : B) \rightarrow \mathbf{Id}(B, f(g b), b)$ as g was chosen more-or-less to make this true by definition. Explicitly, such a function is given by the composite $\mathbf{snd} \circ \text{fst} \circ p$.

It is only slightly more work to define a function $\alpha : (a : A) \rightarrow \mathbf{Id}(A, g(f a), a)$ and it is here that we require that there is exactly one element of each fiber:

$$\alpha : (a : A) \rightarrow \mathbf{Id}(A, g(f a), a)$$

$$\alpha a = \text{cong } (\lambda f \rightarrow \text{fst}(f)) (\text{snd } (p(f a)) (a, \text{refl}))$$

Finally, we must show that $\text{isEquiv}(f)$ is a proposition. To this end, we note that propositions under closed dependent products (Exercise 5.1) and so the result is a direct consequence of Lemma 5.2.7 \square

Lemma 5.2.8. *There is a unique map $\text{idtoequiv} : \mathbf{Id}(\mathbf{U}, A, B) \rightarrow A \simeq B$ such that there is an identification between $\text{fst} \circ \text{idtoequiv}$ and transp .*

Proof. We begin by defining $\text{idtoequiv}(p)$ and prove that it is unique afterwards. Using the identity elimination rule, it suffices to define $\text{idtoequiv}(p)$ when $p = \text{refl}$. That is, we may assume that $A = B$ and show that $\text{isEquiv}(\text{transp } \text{refl})$. We note that the following chain of equalities follows by definition:

$$\text{isEquiv}(\text{transp } \text{refl}) = \text{isEquiv}(\text{subst id } \text{refl}) = \text{isEquiv}(\text{id})$$

To establish $\text{isEquiv}(\text{id})$, we note that it suffices to show $\text{hasTwoSidedInverse}(\text{id})$ which is established by Exercise 5.4.

We now turn to showing that idtoequiv is unique. Suppose we are given another map $f : \mathbf{Id}(\mathbf{U}, A, B) \rightarrow A \simeq B$ along with another identification between $\text{fst} \circ f$ and transp . We wish to construct an element of $\mathbf{Id}(\mathbf{Id}(\mathbf{U}, A, B) \rightarrow A \simeq B, \text{idtoequiv}, f)$. By function extensionality, it suffices to fix $p : \mathbf{Id}(\mathbf{U}, A, B)$ and exhibit an identification $\mathbf{Id}(A \simeq B, \text{idtoequiv } p, f p)$.

We note that $A \simeq B$ is a dependent sum type and so by Exercise 5.6 it suffices to exhibit a pair of identifications:

1. $i : \mathbf{Id}(A \rightarrow B, \text{fst}(\text{idtoequiv } p), \text{fst}(f p))$,
2. an identification in $\text{isEquiv}(\text{fst}(f p))$ between $\text{subst isEquiv } p (\text{snd } (\text{idtoequiv } p))$ and $\text{snd}(f p)$.

We construct i using trans and sym ; we know that both $\text{fst}(\text{idtoequiv } p)$ and $\text{fst}(f p)$ can be identified with $\text{transp } p$. The second identification is *automatic*. By construction, we know that any two elements of $\text{isEquiv}(\dots)$ can be identified. \square

Exercise 5.6. Given $x, y : \sum_{a:A} B a$, construct a function of the following type:

$$(p : \mathbf{Id}(A, \text{fst } x, \text{fst } y)) \rightarrow \mathbf{Id}(B y, \text{subst } B p (\text{snd } x), \text{snd } y) \rightarrow \mathbf{Id}(\sum_{a:A} B a, x, y)$$

Definition 5.2.9. The univalence axiom *Univalence* is defined be $\text{isEquiv}(\text{idtoequiv})$.

It is important to note that univalence is fundamentally a statement *about a universe*. It says something about the behavior of `subst id` which is, in turn, defined with respect to a particular universe. Our informal notation has hidden the presence of the universe slightly; we have erased the subscripts signifying levels and suppressed $\mathbf{El}(-)$. More precisely, we should have written $\mathbf{Univalence}_i$ to signify that the map $\mathbf{idtoequiv}_i : \mathbf{Id}(\mathbf{U}_i, A, B) \rightarrow \sum_{f:\mathbf{El}(A)\rightarrow\mathbf{El}(B)} \mathbf{isEquiv}_i(f)$ is an equivalence. We shall continue to avoid this level of formality for most of this section, but bring it up for the moment to emphasize that core homotopy type theory assumes univalence for every universe \mathbf{U}_i :

Definition 5.2.10. We shall refer to intensional type theory extended with `funext` : `Funext` and $\mathbf{univalence}_i$: $\mathbf{Univalence}_i$ for each universe level i as *core homotopy type theory* (core HoTT).

The definition of univalence may be intimidating at first glance: it involves universes alongside the (surprisingly subtle) `isEquiv` used in two different places. It can be helpful to consider the (equivalent) axiom `hasTwoSidedInverse(idtoequiv)` which can then be recast into three pieces of data:

1. A map $\mathbf{ua} : A \simeq B \rightarrow \mathbf{Id}(\mathbf{U}, A, B)$
2. For each $f : A \simeq B$, an identification $\mathbf{uaBeta} : \mathbf{Id}(A \rightarrow B, \mathbf{transp}(\mathbf{ua} f), \mathbf{fst} f)$
3. for each $p : \mathbf{Id}(\mathbf{U}, A, B)$, an identification $\mathbf{uaEta} : \mathbf{Id}(\mathbf{Id}(\mathbf{U}, A, B), p, \mathbf{ua}(\mathbf{idtoequiv} p))$

Surprisingly, the third piece of data is actually derivable from the first two [Lic16]. Consequently, full univalence can be thought of as the (comparatively simple) assertion that any $f : A \simeq B$ induces an identification $\mathbf{ua} f : \mathbf{Id}(\mathbf{U}, A, B)$ such that $\mathbf{transp}(\mathbf{ua} f) = \mathbf{fst} f$.

While we will not attempt to argue for its consistency, univalence does not contradict intensional type theory. Together with the normalization result for intensional type theory (Theorem 4.2.4), we conclude the following:

Theorem 5.2.11. *Core HoTT is consistent and enjoys normalization and decidable type-checking.*

HoTT does not, however, satisfy canonicity. The quest to resolve this issue will lead us to cubical type theory in Section 5.4.

5.2.2 Homotopy levels

We now take our first steps in homotopy type theory by arguing that univalence actually contradicts $\mathbf{U}(\mathbf{IP})^n$ for all n . We shall see that need not look far to find counterexample of $\mathbf{U}(\mathbf{IP})^n$ in core HoTT: the universe refutes each of these truncation axioms.

In order to formulate this properly, it is convenient to isolate a predicate $\text{hasHLevel} : \text{Nat} \rightarrow \mathbf{U} \rightarrow \mathbf{U}$ in type theory which essentially encodes “ $\text{U}(\text{IP})^n$ holds for A ”. It turns out that this is connected to an important notion of homotopy theory—that of truncation and n -types—and so we shall “types satisfying $\text{U}(\text{IP})^n$ ” as types of “h(omotopy)-level n ”.

An unfortunate fact is that the indexing does not match between h-level n and $\text{U}(\text{IP})^n$ and so, for instance, $\text{hasHLevel } 1 A$ does not signify that A satisfies UIP. The correspondence is that $\text{hasHLevel } (1 + n) A$ holds if A satisfies $\text{U}(\text{IP})^n$. This indexing scheme is standard in homotopy type theory and, unpleasantly, nowhere else. The root of the issue is that there are natural definitions of (-1) -level and (-2) -level types and we wish to capture these in hasHLevel . However, the natural numbers stubbornly start at 0. Accordingly, we must shift everything up by at least 2 when encoding things in type theory. The other mismatch is that a type satisfying UIP behaves like a 0-level type rather than a 1-level type.¹

Up to these reindexing complications, we note that we have already explicitly written out $\text{hasHLevel } 2$ and $\text{hasHLevel } 3$ in Section 4.3 as UIP and UIPIP:

$$\begin{aligned} \text{hasHLevel } 2 A &= (a b : A)(p q : \mathbf{Id}(A, a, b)) \rightarrow \mathbf{Id}(\mathbf{Id}(A, a, b), p, q) \\ \text{hasHLevel } 3 A &= (a b : A)(p q : \mathbf{Id}(A, a, b))(\alpha \beta : \mathbf{Id}(\mathbf{Id}(A, a, b), p, q)) \\ &\rightarrow \mathbf{Id}(\mathbf{Id}(\mathbf{Id}(A, a, b), p, q), \alpha, \beta) \end{aligned}$$

Unfortunately, it is clear that we cannot continue in this way. We must find some way to define $\text{hasHLevel } (\text{suc}(n)) A$ in terms of $\text{hasHLevel } n A$. To this end, let us note that $\text{hasHLevel } 3 A = (a b : A) \rightarrow \text{hasHLevel } 2 (\mathbf{Id}(A, a, b))$. This forms the basis of an inductive definition, but we must figure out what the base case ought to be. Consider $\text{hasHLevel } 3 A$ which ought to encode UIP. By definition, we have now have the following:

$$\begin{aligned} &A \text{ satisfies UIP} \\ &= \text{hasHLevel } 2 A \\ &= (a b : A) \rightarrow \text{hasHLevel } 1 (\mathbf{Id}(A, a, b)) \\ &= (a b : A)(p q : \mathbf{Id}(A, a, b)) \rightarrow \text{hasHLevel } 0 (\mathbf{Id}(\mathbf{Id}(A, a, b), p, q)) \end{aligned}$$

Accordingly, we must choose $\text{hasHLevel } 0 (\mathbf{Id}(\mathbf{Id}(A, a, b), p, q))$ such that this last type is equivalent to A satisfying UIP. Moreover, our goal was that $\text{hasHLevel } n$ should be a predicate on types and so we also wish to ensure that $\text{hasHLevel } 0 A$ is a proposition.

In the final type listed above, if A satisfies UIP then surely p and q are supposed to be equal (as identifications between the same two elements of a type satisfying UIP). At the very least we must have $\text{hasHLevel } 0 X \rightarrow X$, but if we simply take this to be the definition of $\text{hasHLevel } 0 X$ then the result will not be a proposition.

¹Roughly, the lack of non-trivial identifications means the type behaves like a set which is a 0-dimensional object.

However, satisfying UIP tells us a little more than there merely existing an identification. If a type satisfies UIP, then we can actually compute this identification. That is, there is (by definition) a function that take p, q and produces an element of $\mathbf{Id}(\mathbf{Id}(A, a, b), p, q)$. A general result in intensional type theory states that if we are able to define a function $f : (x y : X) \rightarrow \mathbf{Id}(X, x, y)$, then every identification between x and y must be identifiable with $f x y$ [UF13, Lemma 3.3.4]. Accordingly, we can define $\text{hasHLevel } 0 X = \text{isContr}(X)$.

In total then, we define hasHLevel as follows:

$$\begin{aligned} \text{hasHLevel } 0 A &= \text{isContr}(A) \\ \text{hasHLevel } (\text{succ}(n)) A &= (a b : A) \rightarrow \text{hasHLevel } n (\mathbf{Id}(A, a, b)) \end{aligned}$$

Finally, by induction along with Lemma 5.2.7 and Exercise 5.1, we can prove that $\text{hasHLevel } n A$ is a proposition.

Exercise 5.7. Give a description of a type A such that $\text{hasHLevel } 1 A$. (Hint: we have already introduced terminology for this class of types.)

The theory of hasHLevel is extremely rich (see Univalent Foundations Program [UF13, Chapter 7] or Rijke [Rij22, Chapter 12]). However, we will content ourselves with using it only to deliver on a promised result: univalence refutes UIP. That is, we give a type A such that $\text{hasHLevel } 2 A \rightarrow \mathbf{Void}$ holds.

Notation 5.2.12. If $\text{hasHLevel } 2 A$, we shall say that A is a *homotopy set* or *hset* (analogous to a homotopy proposition).

Theorem 5.2.13. $\text{hasHLevel } 2 \mathbf{U} \rightarrow \mathbf{Void}$

Proof. Let us assume $\phi : \text{hasHLevel } 2 \mathbf{U}$ and set about constructing an element of \mathbf{Void} . To this end, we note that we may use univalence to define a function of the following type:

$$\text{ua} : (\mathbf{Bool} \simeq \mathbf{Bool}) \rightarrow \mathbf{Id}(\mathbf{U}, \mathbf{Bool}, \mathbf{Bool})$$

In particular, we obtain a pair of paths $p = \text{ua id}$ and $q = \text{ua not}$ (we leave it to the reader to prove isEquiv not holds). Using ϕ , we then obtain the following:

$$\alpha = \text{fst}(\phi \mathbf{Bool} \mathbf{Bool} p q) : \mathbf{Id}(\mathbf{Id}(\mathbf{U}, \mathbf{Bool}, \mathbf{Bool}), p, q)$$

Again by univalence—specifically uaBeta —we have a further pair of identifications:

$$\mathbf{Id}(\mathbf{Bool} \rightarrow \mathbf{Bool}, \text{transp } p, \text{id}) \quad \mathbf{Id}(\mathbf{Bool} \rightarrow \mathbf{Bool}, \text{transp } q, \text{not})$$

By congruence applied α along with transitivity, we therefore obtain an element $\text{uhoh} : \mathbf{Id}(\mathbf{Bool} \rightarrow \mathbf{Bool}, \text{id}, \text{not})$. Finally, $\text{cong } (\lambda f \rightarrow f \mathbf{true}) \text{uhoh} : \mathbf{Id}(\mathbf{Bool}, \mathbf{true}, \mathbf{false})$ and from here it is straightforward to produce an element of \mathbf{Void} . \square

Corollary 5.2.14. *The type $\mathbf{HSet} = \sum_{A:U} \text{hasHLevel } 2 A$ is not an *hset*.*

In fact, we can to some extent bootstrap this process. Essentially U_0 does not satisfy $\text{hasHLevel } 2$ because it contains a type \mathbf{Bool} which does not satisfy $\text{hasHLevel } 1$. A more complex but morally similar argument tells us that U_{i+1} will not satisfy $\text{hasHLevel } (\text{succ}^{i+3}(0))$ since it contains U_i which does not satisfy $\text{hasHLevel } (\text{succ}^{i+2}(0))$.

Theorem 5.2.15 (Kraus and Sattler [KS15]). *For every external natural number i there is an term of type $\text{hasHLevel } (\text{succ}^{i+2}(0)) U_i \rightarrow \mathbf{Void}$*

Corollary 5.2.16. *Core HoTT does not satisfy $U(\mathbf{IP})^n$ for any n .*

It is natural to wonder if $\text{hasHLevel } 3 U_0$ holds in core HoTT. Phrased differently, did we need to work up the ladder of universes to refute $U(\mathbf{IP})^n$ or could we have refuted them all using a single universe. In fact, the following result is independent of core HoTT:

Proposition 5.2.17. *There exist a type A such that $(n : \mathbf{Nat}) \rightarrow \text{hasHLevel } n A \rightarrow \mathbf{Void}$.*

This type is inhabited in the standard model of HoTT in simplicial sets [KL21] morally because the interpretation of U_0 in this model contains many spaces which are small, but homotopically complex. More generally, it is natural to wonder whether the only source of types not satisfying $\text{hasHLevel } 2$ is the universe.

In fact, in core homotopy type theory this is the case! Every other constructor $\mathbf{\Pi}$, $\mathbf{\Sigma}$, \mathbf{Bool} , *etc.* either has an h-level lower than 2 (*i.e.*, satisfies UIP) or “preserves” types of level k . For instance, the h-level of $\mathbf{\Sigma}(A, B)$ depends on the h-levels of A and B , but $\mathbf{\Sigma}$ does not raise the h-level: if A and B are both of h-level k then so too is $\mathbf{\Sigma}(A, B)$. There is some interesting behavior even with core types: for instance, it turns out that the h-level of $\mathbf{\Pi}(A, B)$ depends only on the h-level of B . However, as it stands the only way to construct a type of h-level higher than 2 is to use universes in some essential way.

Given that we must cope with the hypothetical presence of types with arbitrarily high h-levels, we now discuss a mechanism for defining types of our own which have h-level above 2. These types shall be realized as *higher inductive types* (HITs). We shall see how they allow us to construct homotopically complex types not stemming from the universe and thereby prove Proposition 5.2.17 internally. Remarkably, HITs are useful for far, far more than merely establishing Proposition 5.2.17 and we shall see that they enable us to capture interesting ideas from both homotopy theory and logic within type theory.

5.2.3 Higher inductive types

We also introduce one of the most common extensions to homotopy type theory: higher inductive types. Briefly, once we acknowledge that types may have non-trivial identifications, it becomes natural to include the identifications as part of the specification of inductive types alongside their generating elements.

Remark 5.2.18. We will follow Univalent Foundations Program [UF13] and Rijke [Rij22] (as well as Section 2.5) and treat HITs in an essentially *ad hoc* manner. In particular, we will not concern ourselves with developing a *schema* for higher inductive types [CH19] and instead describe specific examples that ought to exist in any sufficiently expressive schema. \diamond

5.2.3.1 The interval

Let us begin with arguably the simplest possible example and describe a variation on **Bool**. Recall that **Bool** was the type generated by two terms: **true** and **false**. This is enforced by the elimination principle which ensures that every type *believes* that **Bool** merely has these two terms. Let us recall that—roughly—the elimination principle ensures that for each $A : \mathbf{Bool} \rightarrow \mathbf{U}$ there is a section to the following map

$$((b : \mathbf{Bool}) \rightarrow A b) \longrightarrow A \mathbf{true} \times A \mathbf{false}$$

Our goal is to describe a new type **I** which, like **Bool**, is generated by two points **0** and **1** *along with an identification* $\mathbf{seg} : \mathbf{Id}(\mathbf{I}, \mathbf{0}, \mathbf{1})$. Just as before if we are given $A : \mathbf{I} \rightarrow \mathbf{U}$, we have a map from sending a term $a : (i : \mathbf{I}) \rightarrow A i$ to the terms $a \mathbf{0} : A \mathbf{0}$ and $a \mathbf{1} = A \mathbf{1}$, but what is to be done with this “freely added identification” \mathbf{seg} ? It will determine an identification between $a \mathbf{0}$ and $a \mathbf{1}$ by congruence:

$$\mathbf{cong} a \mathbf{seg} : \mathbf{Id}(A \mathbf{1}, \mathbf{subst} A \mathbf{seg} (a \mathbf{0}), a \mathbf{1})$$

If we stitch this all together, we obtain the following map:

$$\mathbf{eval} : ((i : \mathbf{I}) \rightarrow A i) \longrightarrow \sum_{a_0 : A \mathbf{0}} \sum_{a_1 : A \mathbf{1}} \mathbf{Id}(A \mathbf{1}, \mathbf{subst} A \mathbf{seg} a_0, a_1)$$

The idea is crystallize the idea that **I** is generated by **0**, **1**, and \mathbf{seg} by requiring that this map have a section. That is, given $A : \mathbf{I} \rightarrow \mathbf{U}$ we wish to be able to define a function of the following type:

$$\mathbf{rec}_{\mathbf{I}} : (\sum_{a_0 : A \mathbf{0}} \sum_{a_1 : A \mathbf{1}} \mathbf{Id}(A \mathbf{1}, \mathbf{subst} A \mathbf{seg} a_0, a_1)) \rightarrow (i : \mathbf{I}) \rightarrow A i$$

Moreover, we also will ask for an identification between $\mathbf{rec}_{\mathbf{I}} \circ \mathbf{eval}$ and \mathbf{id} . This identification will then morally serve as the “ β equalities” for the elimination principle but up to \mathbf{Id} rather than definitional equality.

This discussion may highlight what we intend to accomplish with higher inductive types like **I** and makes clear what “freely adding identifications” ought to accomplish, but what rules should be added to core HoTT to realize these terms? The answer is, unfortunately, somewhat murky. An often underappreciated appeal of cubical type theory (Section 5.4) is its comparatively clean account of higher inductive types. In core HoTT, it is

not even clear how many definitional equalities may be imposed on rec_I if we were to add it as a primitive term-former. It is possible to imagine the appropriate definitional equations to impose for $\mathbf{0}$ and $\mathbf{1}$, but what about **seg**? If we wish to ensure that the eliminator “computes” on **seg**, we would be forced to specify $\text{cong}(\text{rec}_I(a_0, a_1, p)) \text{seg} = p$.

For this book, we will take the simplest (though least practical) approach and require no definitional equalities. That is, we shall simply postulate rec_I along with its attendant β -equality up to **Id**. While this is unpleasant for performing actual constructions with HITs, it has the advantage of conceptual uniformity and, of course, it will be satisfied by any method of adding HITs to core HoTT.

We will content ourselves with proving only one fact about **I**: we need not have bothered to introduce it as it is equivalent to **Unit**.

Lemma 5.2.19. *The unique map $\mathbf{I} \rightarrow \mathbf{Unit}$ is an equivalence.*

Proof. There are a variety of ways to prove this fact, but we will proceed explicitly. Let us denote the unique map $\mathbf{I} \rightarrow \mathbf{Unit}$ by f and write $g \star = \mathbf{0}$. It is trivial to obtain an identification between $f \circ g$ and id , so we focus on constructing an identification between $g \circ f : \mathbf{I} \rightarrow \mathbf{I}$ and id . By function extensionality, it suffices to construct an identification $\phi : (i : \mathbf{I}) \rightarrow \mathbf{Id}(\mathbf{I}, \mathbf{0}, i)$. For this, we turn to rec_I .

$$\begin{aligned} \alpha &: \mathbf{Id}(\mathbf{Id}(\mathbf{I}, \mathbf{0}, \mathbf{1}), \text{transp } \mathbf{Id}(\mathbf{I}, \mathbf{0}, -) \text{ seg refl, seg}) \\ \alpha &= j(\lambda i_0 i_1 p \rightarrow \mathbf{Id}(\mathbf{Id}(\mathbf{I}, i_0, i_1), \text{transp } \mathbf{Id}(\mathbf{I}, i_0, -) p \text{ refl, } p)) \\ &\quad (\lambda i \rightarrow \text{refl}) \\ &\quad \text{seg} \end{aligned}$$

$$\phi = \text{rec}_I(\text{refl, seg, } \alpha)$$

We emphasize an important point in the above construction: even though **seg** is added as a generating identification in **I**, we may still apply **J** to **seg**. In particular, the addition of generating identifications does not refute our earlier claim that “every type believes that **Id** is generated by **refl**”. \square

5.2.3.2 Suspensions

While **I** was helpful for illustrating the basic ideas behind HITs, it is clearly not very interesting as a type in its own right. Let us therefore turn to a family of HITs which were not already definable in type theory: suspensions ΣA .

We begin by specifying the generators for ΣA . As before, we shall work with these additional axioms in our type theory alongside an axiom $\Sigma : \mathbf{U} \rightarrow \mathbf{U}$ but without any definitional equalities. There are two generating terms **north**, **south** : ΣA and a generating identification **merid** $a : \mathbf{Id}(\Sigma A, \text{south}, \text{north})$ for each $a : A$.

Remark 5.2.20. It may be helpful to note that this generalizes **I** by allowing for not just a single generating identification between a pair of elements but any number we choose. \diamond

Remark 5.2.21. We shall see in Section 5.2.4 that ΣA can be visualized as effectively as some sort of globe with a north and south pole joined by a fresh meridian for each $a : A$. We note that there is no relation between dependent sum types and suspension types other than both starting with an “s” and therefore being traditionally written with a Σ . \diamond

As before, we use these to define an evaluation map for a type $B : \Sigma A \rightarrow \mathbf{U}$:

$$\begin{aligned} \text{eval} : ((s : \Sigma A) \rightarrow B s) \\ \rightarrow \sum_{b_s : B \text{ south}} \sum_{b_n : B \text{ north}} (a : A) \rightarrow \mathbf{Id}(B \text{ north}, \text{subst } B (\text{merid } a) a_s, a_n) \end{aligned}$$

Finally, we axiomatize the elimination principle as a section to this map:

$$\begin{aligned} \text{rec}_\Sigma : (\sum_{b_s : B \text{ south}} \sum_{b_n : B \text{ north}} (a : A) \rightarrow \mathbf{Id}(B \text{ north}, \text{subst } B (\text{merid } a) a_s, a_n)) \\ \rightarrow (s : \Sigma A) \rightarrow B s \end{aligned}$$

The importance to ΣA is that it can be made to refute $\mathbf{U}(\mathbf{IP})^n$ for any n by carefully choosing A . While a proof is out-of-scope for our purposes, it is possible to show that $\Sigma^n \mathbf{Bool}$ does not satisfy $\text{hasHLevel}(1 + n)$ given $n : \mathbf{Nat}$.²

Theorem 5.2.22 (Theorem 8.6.17 [UF13]). *For any $n : \mathbf{Nat}$, the suspension $\Sigma^n \mathbf{Bool}$ does not satisfy $\text{hasHLevel}(1 + n)$.*

Corollary 5.2.23. *The type $X = \sum_{n : \mathbf{Nat}} \Sigma^n \mathbf{Bool}$ does not satisfy $\text{hasHLevel } n$ for any $n : \mathbf{Nat}$.*

Proof. We observe that $\Sigma^n \mathbf{Bool}$ is a *retract* of X for any $n : \mathbf{Nat}$; there are maps $e : \Sigma^n \mathbf{Bool} \rightarrow X$ and $r : X \rightarrow \Sigma^n \mathbf{Bool}$ such that $r \circ e$ is the identity. One can argue that if A is a retract of B and $\text{hasHLevel } n B$ then $\text{hasHLevel } n A$. However, we know that $\text{hasHLevel}(1 + n)(\Sigma^n \mathbf{Bool})$ does not hold and so X must not satisfy $\text{hasHLevel } n$ for any n . \square

5.2.3.3 Set truncations

For our final example of higher inductive types, we turn to an example which is truly inductive. For both **I** and Σ , we postulated types with constructors generating both new elements and identifications, but in both cases none of these constructors took an element of either **I** or Σ . This is akin to **Bool** or **Void** and in contrast to a type like **Nat** where

²While it is much harder to prove, it is also the case that $\text{hasHLevel } n(\Sigma^2 \mathbf{Bool})$ is refuted by at least some models of homotopy type theory for all $n > 4$. To the authors’ knowledge, it is still open whether this classic result due to Serre [Ser53] holds directly within HoTT.

$\text{succ}(-)$ takes an element of Nat . Our last example of set truncation's, $|A|$, is more the cast of Nat .

First, $|A|$ is generated by the following constructors:

- If $a : A$ then $[a] : |A|$
- If $p, q : \text{Id}(|A|, a, b)$ then $\text{trunc} : \text{Id}(\text{Id}(|A|, a, b), p, q)$

In summary, there is a map $A \rightarrow |A|$ and we further ensure that $|A|$ has h-level 2: every pair of identity proofs can be identified. The second constructor ranges over elements of $|A|$ and gives $|A|$ its “inductive” character. The force of this last constructor is in the following observation:

Lemma 5.2.24. $\text{hasHLevel } 2 \ |A|$.

Proof. Unfolding, we see that a witness for this type is *precisely* trunc . \square

We shall see that $|A|$ occupies a particular place among types satisfying $\text{hasHLevel } 2$; it is the “best approximation to A satisfying this predicate”. In other words, $| - |$ allows us to replace a type with one satisfying UIP. To make this statement precise and prove it, we must specify the elimination principle for $|A|$.

An interlude on displayed algebra Before we arrive at the elimination principle, we must specify what it means to be a $|A|$ -display algebra just as we did for Nat in Section 2.5.3. We can follow the same process there and mechanically derive the following (rather complex) definition of what it is required for $B : |A| \rightarrow \mathbf{U}$ to be a displayed algebra over $|A|$ when it comes equipped with the following:

- A function $\overline{[-]} : (a : A) \rightarrow B([a])$.
- Given $p, q : \text{Id}(A, a_0, a_1)$ along with $b_0 : B(a_0), b_1 : B(a_1), \bar{p} : \text{Id}(B(a_1), \text{subst } B p b_0, b_1)$, and $\bar{q} : \text{Id}(B(a_1), \text{subst } B q b_0, b_1)$, an element of the following type:

$$\overline{\text{trunc}} : \text{Id}(\text{Id}(B(a_1), \text{subst } B q b_0, b_1), \text{subst}(\lambda r \rightarrow \text{Id}(B(a_1), \text{subst } B r b_0, b_1)) \text{trunc } \bar{p}, \bar{q})$$

This specification expresses some of the complexity that can come from working with \mathbf{J} and intensional identity types, especially in the context of HoTT. The need to apply “dependent substitutions” quickly leads to impossible to read types. Fortunately, in this particular case we can give a much simpler specification:

Lemma 5.2.25 (Lemma 6.9.1 [UF13]). *The data required for $B : |A| \rightarrow \mathbf{U}$ to be a displayed $|A|$ -algebra is equivalent to a function $\overline{[-]} : (a : A) \rightarrow B([a])$ along with $(a : A) \rightarrow \text{hasHLevel } 2 \ (B a)$.*

We can package things up even further by recalling the definition $\mathbf{HSet} = \sum_{A:U} \mathbf{hasHLevel} 2 A$. We may then omit the final requirement of the above lemma by instead requiring $B : |A| \rightarrow \mathbf{HSet}$.

With this digression, we may state the elimination principle for $|A|$ as a section to the following canonical map for any $B : |A| \rightarrow \mathbf{HSet}$ given by precomposing with $\overline{[-]}$:

$$((a : |A|) \rightarrow B(a)) \rightarrow ((a : A) \rightarrow B([a]))$$

Explicitly, we have the following elimination rule:

$$\mathbf{rec}_{|A|} : ((a : A) \rightarrow B([a])) \rightarrow (a : |A|) \rightarrow B(a)$$

Finally, we make good on our earlier promise:

Theorem 5.2.26. *Given any $B : \mathbf{HSet}$, the type $A \rightarrow B$ is equivalent to $|A| \rightarrow B$.*

Proof. This is an immediate consequence of the elimination principle specialized to a non-dependent motive. \square

This theorem shows that “from the perspective of an h-set”, $|A|$ is equivalent to A . In fact, one can define a version of $|A|$ for every possible h-level. These allow one to better and better approximate A by a sequence of types which satisfy some version of $\mathbf{U}(\mathbf{IP})^n$. In practice, $|A|$ and its cousin for homotopy propositions arise most frequently.

5.2.4 A sampling of homotopy type theory

This section has been a whirlwind tour of some of the most interesting type-theoretic aspects of homotopy type theory. It is, however, an impossible task to fully discuss HoTT in merely a section of this book; multiple books have been written on the topic. Instead, we content ourselves by concluding our excursion through HoTT with a discussion of a few applications of the theory and pointers to the relevant literature. We (somewhat artificially) divide these applications into three separate tracks: synthetic homotopy theory, quotients and colimits in type theory, and the structure-identity principle.

Remark 5.2.27. We once more remind the reader that the purpose of this book is to help them engage with the literature. They are therefore encouraged to do the legwork of studying some of the cited papers and books. \diamond

Synthetic homotopy theory As one might have inferred from the name, HoTT is closely related to homotopy theory. We have thus far avoided any real discussion of the field, but homotopy theory arose from the study of *(algebraic) topology*: a subfield of geometry concerned with study shapes up to twisting, stretching, and compression.

While homotopy theory has since spread out to touch many distinct areas of mathematics (algebraic geometry, representation theory, mathematical physics, and number theory, among many others!) many of the classical results of the field concern properties of familiar shapes such as circles, spheres, and tori.

Interestingly, the connection between HoTT and homotopy theory allows us to recast these shapes as particular types and thereby reproduce classical results from algebraic topology purely within type theory. The resulting theorems *synthetic* theorems apply in greater generality (to any model of HoTT) and often yield shorter and more conceptual proofs. Some of this is introduced in Univalent Foundations Program [UF13] and Rijke [Rij22].

For instance, one can define a type which behaves like a circle: $S^1 = \Sigma \mathbf{Bool}$. Pictorially, this is a type with 2 elements joined by 2 identifications. If we view elements as points and identifications as paths, this is a valid description of a circle. A classical result in algebraic topology is to fully characterize the maps $S^1 \rightarrow S^1$ which send **south** to **south** (so-called *pointed* maps). Within the framework of homotopy type theory, we may characterize this into something more familiar to type theorists:

Lemma 5.2.28. *The type of pointed maps $S^1 \rightarrow_* S^1$ is equivalent to $\mathbf{Id}(S^1, \mathbf{south}, \mathbf{south})$.*

One of the classical results of homotopy theory (and homotopy type theory) is a complete characterization of this type. One approach to this result is through careful analysis of type families $B : S^1 \rightarrow \mathbf{U}$ —the HoTT version of *covering spaces*. In particular, note that by construction B consists of (1) a type $A : \mathbf{U}$ and (2) an identification/equivalence $e : A \simeq A$. For instance, one can define B via \mathbf{Bool} and \mathbf{not} . One can similarly consider a type with three elements and the isomorphism corresponding to successor modulo three. In the limit, we find the universal type family to be defined from \mathbb{Z} and the \mathbf{succ} from which Licata and Shulman [LS13] concluded the following:

Theorem 5.2.29 (Corollary 8.1.10 [UF13]). $\mathbf{Id}(S^1, \mathbf{south}, \mathbf{south}) \simeq \mathbb{Z}$.

While this is one of the earliest results of synthetic homotopy theory, many classical results have now been defined others are being actively studied. We refer the reader to Univalent Foundations Program [UF13] and Rijke [Rij22] for textbook discussions of synthetic homotopy theory. We also highlight the theses Van Doorn [VD18] and Brunerie [Bru16] for the construction of the Serre spectral sequence and the computation of $\pi_4(S^3)$ respectively. A particularly notable instance of synthetic homotopy theory is the proof of the Blakers-Massey theorem due to Hou (Favonia) et al. [Fav+16]. This synthetic proof is significantly simpler than existing “analytic” proofs and has since led to various generalizations in homotopy theory.

Descent and colimits A surprising fact about HoTT is that many “quotient”-type constructions when realized as HITs behave much better than attempts to codify them in extensional type theory.³ Formally, one may say that all quotient-type constructions or *colimits* in HoTT exhibit *effectivity* or *descent* [Lur09; Rez10; AJ21].

To see the utility of descent in action, let us consider how one might realize the quotient of $A : \mathbf{HSet}$ by an equivalence relation $R : A \rightarrow A \rightarrow \mathbf{HProp}$ as a HIT A/R . This type has three constructors:

- For each $a : A$, there is an element $[a] : A/R$.
- For each $r : R a_0 a_1$, there is an identification $[r] : \mathbf{Id}(A/R, [a_0], [a_1])$
- For each $p, q : \mathbf{Id}(A/R, x, y)$, there is an identification $\mathbf{Id}(\mathbf{Id}(A/R, x, y), p, q)$

Remark 5.2.30. The final constructor ensures that A/R is an hset. Without this modification, we may accidentally encounter e.g. $S^1 = \mathbf{Unit}/(\lambda_ \rightarrow \mathbf{Unit})$. \diamond

We will omit the elimination principle, but invite the reader to attempt to devise what it should be themselves and consult Univalent Foundations Program [UF13, Lemma 6.10.3] to check their answer.

It is clear that if a_0 and a_1 are related by R , then $[a_0]$ and $[a_1]$ are identified. What is less clear, however, is that this condition is both sufficient and necessary. This is one of the most basic expressions of descent and often referred to as the *effectivity* of quotients:

Theorem 5.2.31 (Lemma 10.1.8 [UF13]). $\mathbf{Id}(A/R, [a_0], [a_1]) \simeq R a_0 a_1$.

This property is crucial for using quotients in ordinary mathematics. Imagine, for instance, attempting to represent \mathbb{Q} as the quotient of $\mathbb{Z} \times \mathbf{Nat}$ and being unable to prove that (a, b) and (c, d) are identified just when $a/b = c/d$.

What is perhaps surprising to the reader is not that this property holds in HoTT, but that it frequently does not hold in intensional type theory! Returning to the discussion in at the beginning of this section, the rough issue is that effectivity of quotients holds only if we consider equivalence relations valued in propositions which are *definitionally* propositional, not merely homotopy propositions. This is something of a quagmire for intensional type theory: in most situations homotopy propositions are preferable *except* for quotient types where strict propositions are superior.

Fortunately, homotopy type theory rectifies this situation by ensuring that homotopy propositions are the correct answer in all situations. This improvement is a direct consequence of allowing additional identifications in the universe *i.e.*, univalence.

We also note that effectivity is usually a somewhat special property of quotients by equivalence relations. Even working in ordinary set theory, even though one can ask

³This is, in fact, a manifestation of the slogan that “colimits work better in ∞ -categories”.

whether any *colimit* construction satisfies effectivity, it will hold in generality only for (1) disjoint unions and (2) quotients by equivalence relations. The remarkable fact of homotopy theory and HoTT is that effectivity holds for *all* colimits. For instance, a more modern perspective on Theorem 5.2.29 is that it is a direct consequence of effectivity for the HIT used to define S^1 .

Advanced Remark 5.2.32. While thoroughly out of scope, it is worth mentioning that effectivity of all colimits is more than just a consequence of univalence, it is *equivalent to it* when working categorically. For more on this perspective, see Lurie [Lur09, Section 6.1.6], Gepner and Kock [GK16], or Rasekh [Ras21]. An approachable talk on the subject was also given by Anel [Ane19]. \diamond

Structure and Identity Principle An inarguable lesson of twentieth century mathematics is that whenever one encounters a structure, the task is to find an appropriate notion of map (*homomorphism*) between instances of this structure. While *appropriate* depends on the situation, a fundamental principle is that any property of interest of *e.g.*, groups is stable under *e.g.*, invertible group homomorphisms (*isomorphisms*). A remarkable fact of univalence is the *structure identity principle* (SIP): for any type of structured objects, univalence ensures that the corresponding identity type is equivalent to isomorphisms.

To see this principle concretely, consider the following type of monoids:

Mon =

-- The carrier

$\sum_{A:\mathbf{HSet}}$

-- The operations

$\sum_{\epsilon:A}$

$\sum_{(\cdot):A \rightarrow A \rightarrow A}$

-- The equations

$\sum_{(a:A) \rightarrow \mathbf{Id}(A, a \cdot \epsilon, a)}$

$\sum_{(a:A) \rightarrow \mathbf{Id}(A, \epsilon \cdot a, a)}$

$((a_0 a_1 a_2 : A) \rightarrow \mathbf{Id}(A, a_0 \cdot (a_1 \cdot a_2), (a_0 \cdot a_1) \cdot a_2))$

It is easy enough to define the type of monoid homomorphisms: a map between carriers which commutes with multiplication and ϵ . What is remarkable is that monoid *isomorphisms* fall out without additional work:

Theorem 5.2.33. *If $A, B : \mathbf{Mon}$ the type $\mathbf{Id}(\mathbf{Mon}, A, B)$ is equivalent to the type of monoid isomorphisms $A \cong B$.*

Remark 5.2.34. In a certain sense, this theorem is a repeated application of Exercise 5.6 and massaging identity types. \diamond

What is most profitable about SIP is its robustness: it applies not only to monoids or even algebraic structures, but to everything from partial orders to Petri nets to dagger categories. A thorough study of this principle is carried out in the book by Ahrens et al. [Ahr+22].

5.3★ *Motivating and justifying univalence*

This chapter is devoted to discussing the consequences of the univalence axiom, but how could someone come up with this axiom in the first place? In Section 5.2, we have seen this myriad ways that univalence can be used to enhance or improve type theory. However, one cannot deny that $\text{Id}(\mathbf{U}, A, B) \simeq (A \simeq B)$ is hardly an intuitive statement; some evidence for this can be found in the fact that it is quite easy to define \simeq incorrectly and thereby produce a provably false axiom. Furthermore, it is quite difficult to construct a semantic model of univalence as it requires some insight and technology from homotopy theory. All of this begs the question: why did Voevodsky and others believe univalence enough to attempt to prove it and how does one go about showing that univalence is consistent?

The goal of this section is to answer these questions. We will motivate how one inevitably arrives at univalence by striving to find a characterization of universes in line with dependent products or dependent sums: to find an isomorphism $\text{Tm}(\Gamma, \mathbf{U})$ which allows us to construct all the relevant structure on \mathbf{U} .

This question will naturally lead us to higher-dimensional analogs of sets (groupoids and their generalizations). We shall then see that if we view type theory as a language for dealing with (higher) groupoids, then univalence is an exceptionally natural principle to require. This particular story is hardly novel.⁴ However, we will attempt to present this line of ideas at the same level of mathematical sophistication as the rest of this book.

Remark 5.3.1. We do not wish to downplay, however, the insight required by Hofmann and Streicher [HS98], Awodey and Warren [AW09], Voevodsky [Voe06], and others to realize that type theory could be such a language! From our perspective, the shock is not so much the univalence axiom itself, but the rich connection between type theory and homotopy theory. \diamond

This section will have relatively little type theory in it. Rather, we shall investigate the behavior of universes in set theory and groupoids and merely sketch how to connect the following constructions to type theory via the semantic interpretations we have seen in Sections 3.5 and 4.3, respectively. While we do not have the space to fully develop the mathematics necessary to model full univalence, we do show how to justify the inclusion of a single univalent universe into type theory.

⁴The role of groupoids, higher groupoids, and eventually homotopy theory in constructing classifying or moduli spaces has been a running theme in the field since near its inception.

5.3.1 Propositions and propositional univalence

Part of the problem with attempting to argue about $\mathbf{Id}(U, A, B)$ is that U is a rather complex structure! We introduced it to manage thorny size issues we encountered in Section 2.6 when attempting to internalize the judgment $\Gamma \vdash A$ type. Despite the clear motivation, it can be rather difficult to conceptualize e.g., what set is used to interpret U in the set model. In order to cut through some of the size issues and focus on the essentials, we shall repeat the beginning of Section 5.2 and discuss the semantics of the universe of propositions (Section 2.8) to avoid Russelian paradoxes. We shall see that in this case everything works as well as can be reasonably expected and that the universe of propositions can be fully axiomatized by a natural isomorphism.

We begin by defining $\Omega = \{\top, \perp\}$. This set then contains a (representative) for every set with at most one element *i.e.*, every proposition. There are no size issues here because there are only two such propositions: “true” \top and “false” \perp .

We shall now show that Ω satisfies a mapping-in universal property akin to those governing Π or Σ . In the same way that a map F from A to the “type of all types” U is a dependent type assigning to each element $a \in A$ a type $F(a)$, a map P from A to the “set of all propositions” Ω is a “dependent proposition” assigning to each $a \in A$ a proposition $P(a)$, more commonly known as a predicate on A . Contemplating further, the reader will notice that a predicate on A is precisely a subset of A which leads to the following:

Theorem 5.3.2. *There is an isomorphism between $A \rightarrow \Omega$ and $\mathcal{P}(A)$ (subsets of A).*

Proof. We first describe procedure for converting a map $f : A \rightarrow \Omega$ into a subset of $\Phi(f) \subseteq A$. We send f to the preimage of $\{\top\}$ or “the subset of A mapped by f to true elements”. Concretely:

$$\Phi(f) = \{a \in A \mid f(a) = \top\}$$

For the reverse direction, we must send a subset $A_0 \subseteq A$ to a map $\chi(A_0) : A \rightarrow \Omega$. We define this as follows:

$$\chi(A_0)(a) = \begin{cases} \top & a \in A_0 \\ \perp & a \notin A_0 \end{cases}$$

We leave it to the reader to check that $\chi(\Phi(f)) = f$ and $\Phi\chi(S) = S$. □

The reader may wonder how our intuition that \top signifies the true proposition and that \perp denotes false factors into the above proof. Inspecting the proof, we see that the subset determined by a map $f : A \rightarrow \Omega$ is exactly those elements of A sent to \top . However, nothing in particular forced this: we could just as well have reversed the roles of \top and \perp in the above proof and still construct a valid (if perhaps confusing) isomorphism.

While we have not yet proven it, the above isomorphism is *natural* in A . We shall see that a natural family isomorphism $\iota : (A \rightarrow \Omega) \cong \mathcal{P}(A)$ both determines and is determined

by a choice of which of the two elements of Ω is intended to denote the true proposition. We have just seen how a choice of which element is true determines an isomorphism. For the reverse direction, we shall see that the subset $\iota(\mathbf{id} : \Omega \rightarrow \Omega) \subseteq \Omega$ will consist of exactly those propositions that ι “believes” are true. Furthermore, we can use the naturality of ι to show that this subset must consist of precisely one element, the true proposition.

In order to crystallize this discussion, we must once again revisit the concept of a natural bijection. We wish to talk about a family of isomorphisms between $A \rightarrow \Omega \cong \mathcal{P}(A)$, natural in A . In this instance, naturality says that the following square commutes for each map of sets $f : A \rightarrow B$:

$$\begin{array}{ccc} B \rightarrow \Omega & \xrightarrow{\quad} & \mathcal{P}(B) \\ f^* \downarrow & & \downarrow f^{-1} \\ A \rightarrow \Omega & \xrightarrow{\quad} & \mathcal{P}(A) \end{array}$$

Here f^* sends a map $g : B \rightarrow \Omega$ to $g \circ f : A \rightarrow \Omega$ and f^{-1} takes a subset $S \subseteq B$ to its inverse image $\{a \in A \mid f(a) \in S\}$ under f .

Theorem 5.3.3.

- An element $p \in \Omega$ determines a natural isomorphism $\iota_A^p : (A \rightarrow \Omega) \cong \mathcal{P}(A)$
- A natural isomorphism $\iota_A : (A \rightarrow \Omega) \cong \mathcal{P}(A)$ determines an element $p^i \in \Omega$

Moreover, these constructions are inverse.

Proof. Given $X_0 \subseteq X$, we have already seen how to define ι_A (we leave it to the reader to check that the isomorphism constructed in Theorem 5.3.2 is natural).

For the reverse direction, we take $X_0 = \iota_X(\mathbf{id})$. We must first argue that $|X_0| = 1$. The naturality condition implies that $|X_0| \neq 0$. Indeed, it were then for every $f : A \rightarrow X$ we would have $\iota_A(f) = \iota_A(\mathbf{id} \circ f) = f^{-1}(\iota_X \mathbf{id}) = \emptyset$ which would contradict our assumption that ι_A surjects onto $\mathcal{P}(A)$. A similar argument shows that $|X_0| \neq 2$. For a map $f : A \rightarrow X$ we would have $\iota_A(f) = \iota_A(\mathbf{id} \circ f) = f^{-1}(\iota_X \mathbf{id}) = A$ and derive the same contradiction.

The calculation that these procedures are inverse offers no surprises and is omitted. \square

Advanced Remark 5.3.4. The above can be summarized by stating that Ω represents \mathcal{P} . The discussion around one element subsets is then an instance of the Yoneda lemma. \diamond

The reader should compare this with Section 2.4: we have the “type” Ω and a natural isomorphism characterizing maps into Ω . This description then isolates what characterizes the universe of propositions. There is no need to worry about how many elements Ω has or which are marked as “true” or not; it suffices to simply insist that there be a natural

isomorphism between $X \rightarrow \Omega$ and $\mathcal{P}(X)$. One often summarizes this by saying that a map into Ω classifies a subset of the domain of the map.

Advanced Remark 5.3.5. In category theory, subsets are generalized to *subobjects* and so category theorists refer to objects behaving like Ω as *subobject classifiers*. We shall adopt this terminology as well. \diamond

While this satisfies our goal of uniquely characterizing Ω , this mapping-in property offers additional advantages. Most importantly, we are able to use it to construct other operations on Ω by appealing to the relevant existing constructions on $\mathcal{P}(-)$. In other words, this natural isomorphism gives a small but very useful API for Ω .

In more detail, we can use this correspondence to define operations on Ω by defining “natural families” of operations on power sets. With it, we can define boolean logic (\wedge , \vee , \neg , \dots) on Ω . This should be contrasted with Section 2.6 where we were required to manually stipulate closure under various connectives to compensate for the lack of either a mapping-in or mapping-out property for $(\mathbf{U}, \mathbf{El})$.

Rather than developing this systematically—a task best done with category theory and the Yoneda lemma—we explain a particular example: the entailment relation $(\Rightarrow) \subseteq \Omega \times \Omega$. In particular, we can use our natural bijection to define \Rightarrow as:

$$p \Rightarrow q = \iota_{\{\star\}}(p) \subseteq \iota_{\{\star\}}(q)$$

In other words, we note that $p, q \in \Omega$ correspond to subsets of the singleton set $\{\star\}$ and repurpose \subseteq on $\mathcal{P}(\{\star\})$ to define \Rightarrow . More generally, we can show that this construction is part of a bijection: \Rightarrow suffices to determine \subseteq and \subseteq (obviously) determines \Rightarrow . All the salient properties of \Rightarrow (or, indeed, any operation on Ω) can be studied by passing back to their effect on subsets. Consequently, we are able to derive the following lemma:

Lemma 5.3.6. *If $p \Rightarrow q$ and $q \Rightarrow p$ then $p = q$.*

Proof. This is a direct consequence of the fact that \subseteq is a partial order. \square

We can use this result as the basis for the consistency of propositional univalence as introduced in Section 5.3.1. Essentially, we must show that we can arrange for (homotopy) propositions in type theory to be realized as elements of Ω . While possible, this is a somewhat delicate task. One must ensure that it is not necessary to equate $\{(\star, \star)\}$ and $\{\star\}$, but these complications can be avoidable. Roughly, we explicitly restrict the interpretation of types from arbitrary (small) sets to (small) cardinals. This choice is somewhat artificial and requires us to explicitly appeal to the axiom of choice in several places, but ensures that the interpretations of e.g., $\mathbf{Unit} \times \mathbf{Unit}$ and \mathbf{Unit} coincide.

Since this is somewhat ad-hoc and specific to sets, one usually formulates propositional univalence somewhat differently than we did in Section 5.3.1 and more inline with Section 2.8. Rather than working with \mathbf{HProp} as defined above, we define a new type Ω along

with a decoding map $[-] : \Omega \rightarrow \mathbf{HProp}$. We require that (1) Ω satisfies propositional univalence and (2) For every $A : \mathbf{HProp}$ there exists an element $\phi : \Omega$ such that $A \cong [\phi]$. In fact, due to (1), the element ϕ in (2) is unique up to \mathbf{Id} .

By extending the ordinary set model to interpret the type Ω as the set Ω described above, we obtain the following result:

Theorem 5.3.7. *Both intensional and extensional type theory extended with a type Ω satisfying (1) and (2) above is consistent.*

5.3.2 An intermezzo on families of sets

Before moving on to our next example of a universe, we must return to the question of representing families of sets. In our discussion of Ω we have dealt with families of propositions—predicates—and these can be encoded as subsets. Presently, however, we will consider families of non-subsingleton sets and subsets no longer suffice. We will instead opt for a generalization and encode families of sets by maps.

Definition 5.3.8. A family of sets over X is simply a map $f : Y \rightarrow X$.

Notation 5.3.9. When a map $f : Y \rightarrow X$ is intended to represent a family, we will write $f : Y \multimap X$ to help visually distinguish it.

Where have the families of sets gone in this definition? They are hidden away as the *preimages* $f^{-1}(x)$. Indeed, if we imagine a collection of sets $(Y_x)_{x \in X}$, we can bundle this together into a single set $Y = \coprod_{x \in X} Y_x$ along with a projection $\pi : Y \multimap X$. In this situation, the preimage (sometimes *fiber*) over x precisely recovers Y_x . More generally, since any function $f : Y \rightarrow X$ induces a bijection $Y \cong \coprod_{x \in X} f^{-1}(x)$, we can see that every $f : Y \multimap X$ determines a collection $(Y_x)_{x \in X}$.

By bundling up sets as a disjoint union and projecting onto X , we recast “a function producing sets” into something purely expressible in terms of sets and functions between them. Since we have many tools for studying and manipulating sets and functions between them, this maneuver will save us some work in what is to come.

Definition 5.3.10. A family $f : Y \multimap X$ is a family of finite sets if each $f^{-1}(x)$ is a finite set.

5.3.3 Finite sets and isomorphism reflection

Continuing now with our main story, let us pick the second simplest universe of sets: the universe of finite sets. Our goal is once more to find a set X for which maps $Y \multimap X$ correspond in some way to finite families over Y . Just as before, we begin by isolating a set which contains a unique code for each finite set and, once again, we do not have to

look far for such a set. This time our universe shall be the set of natural numbers \mathbb{N} and we shall think of $n \in \mathbb{N}$ as the code for a set with n elements.

To state the mapping-in property of \mathbb{N} , we will have to be a bit more precise about how a map $Y \rightarrow \mathbb{N}$ *corresponds* to a family of finite sets. In particular, we shall see that such a map does not precisely determine a finite family, but rather a family *up to equivalence*:

Definition 5.3.11. Two families $f : Y \rightarrow X$ and $g : Z \rightarrow X$ are *equivalent* if there is a bijection $\alpha : Y \cong Z$ such that $f = g \circ \alpha$.

Definition 5.3.12. We write $\mathbf{Fam}_{\text{fin}}(X)$ for the set of finite families over X taken up to equivalence.⁵

Remark 5.3.13. To connect this with subsets, we note that a subset of X corresponds precisely to an equivalence class of families $f : Y \rightarrow X$ such that each $f^{-1}(x)$ contains *at most one element*. That is, a subset of X corresponds to a family of subsingleton sets. \diamond

Exercise 5.8. Construct a bijection between the subset of $\mathbf{Fam}_{\text{fin}}(X)$ spanned by families of subsingleton sets and $\mathcal{P}(X)$.

We now arrive at the mapping-in property for \mathbb{N} :

Theorem 5.3.14. *There is a natural bijection between maps $X \rightarrow \mathbb{N}$ and $\mathbf{Fam}_{\text{fin}}(X)$.*

Before proving this result, it is helpful to construct the equivalent of the subset $\{\top\} \subseteq \Omega$ for \mathbb{N} . That is, we define a finite family over \mathbb{N} which decodes n to such a finite set with n elements:

$$\begin{aligned} \pi : (\coprod_{n:\mathbb{N}} \{m \in \mathbb{N} \mid m < n\}) &\rightarrow \mathbb{N} \\ \pi(n, _) &= n \end{aligned}$$

We will write \mathbb{N}^\bullet for the domain of π , but the centerpiece of this definition is not the domain but the size of the preimages $\pi^{-1}(n)$.

Proof of Theorem 5.3.14. We first construct a map $\Phi : (X \rightarrow \mathbb{N}) \rightarrow \mathbf{Fam}_{\text{fin}}(X)$. Given a map $f : X \rightarrow \mathbb{N}$, we define the family $\Phi(f) = f^* \pi$ to be the following family over X :

$$f^* \pi : f^* \mathbb{N}^\bullet = \{(x \in X, n^\bullet \in \mathbb{N}^\bullet) \mid f(x) = \pi(n^\bullet)\} \rightarrow X \quad f^* \pi(x, _) = x$$

For the reverse, we send a family of finite sets $\sigma : X^\bullet \rightarrow X$ to the following map:

$$\Phi^{-1}(\sigma)(x) = |\sigma^{-1}(x)|$$

⁵A priori, this is a proper class but we shall see momentarily that it is, in fact, actually a set.

By assumption, σ is a family of finite sets and so $\sigma^{-1}(x)$ is a natural number for all x . We leave it to the reader to check that Φ and Φ^{-1} are inverses.

To show that Φ is natural, fix a map $f : X \rightarrow X'$. We must first specify the a map $f^* : \mathbf{Fam}_{\text{fin}}(X') \rightarrow \mathbf{Fam}_{\text{fin}}(X)$ and then show that the following square commutes:

$$\begin{array}{ccc} B \rightarrow \mathbb{N} & \longrightarrow & \mathbf{Fam}_{\text{fin}}(B) \\ f^* \downarrow & & \downarrow f^* \\ A \rightarrow \mathbb{N} & \longrightarrow & \mathbf{Fam}_{\text{fin}}(A) \end{array}$$

The map $f^* : \mathbf{Fam}_{\text{fin}}(X') \rightarrow \mathbf{Fam}_{\text{fin}}(X)$ is an instance of a pullback and a generalization of the definition of $f^*\pi$ we used to define Φ . Given a family $\sigma : X^\bullet \rightarrow X'$ we define $f^*\sigma$ to be the family $f^*X^\bullet \rightarrow X$ given by the following construction:

$$f^*X^\bullet = \{(x \in X, x^\bullet \in X^\bullet) \mid f(x) = \sigma(x^\bullet)\} \quad f^*\sigma(x, _) = x$$

We leave the verification that the above diagram commutes as an exercise. \square

Just as in Section 5.3.1, we can sharpen this result: the choice of (natural) isomorphisms $(X \rightarrow \mathbb{N}) \cong \mathbf{Fam}_{\text{fin}}(X)$ precisely determines $\pi : \mathbb{N}^\bullet \rightarrow \mathbb{N}$ *up to equivalence*. Let us note that \mathbb{N} satisfies a property which almost looks like (propositional) univalence:

Corollary 5.3.15 (Equivalence implies equality). *If a pair of families σ_0, σ_1 are equivalent over X , then they induce equal maps $X \rightarrow \mathbb{N}$.*

While we will not go into details, one can parlay this into the following result:

Theorem 5.3.16 (Theorem 8.3.2 [Win20]). *Extensional type theory is consistent with the following axiom: $(A B : \mathbb{U}) \rightarrow A \cong B \rightarrow \mathbf{Eq}(\mathbb{U}, A, B)$.*

This is close to univalence: equivalent families induce equal elements of the universe. However, we do not have a bijection between equivalences $\sigma_0 \cong \sigma_1$ and equalities of the resultant maps $X \rightarrow \mathbb{N}$; maps can be equal in at most one way but families can be equivalent in multiple distinct ways. This mismatch is why the mapping-in property for \mathbb{N} must work with families up to equivalence and we shall see that this makes it far less useful than the property governing Ω .

The root of this deviation is that the passage from a family π , to a function to \mathbb{N} , back to a family preserves only the cardinality of the fibers of π and necessarily discards the actual precise finite sets which were used. As a result, however, if we wish to define an operation on \mathbb{N} using a corresponding operation on finite families, we must ensure that equivalent families are sent to *identical* results. This is not an issue if the target of such an operation is $\mathbf{Fam}_{\text{fin}}(-)$ e.g., if we define $(\rightarrow) : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$, but causes issues for a number of other basic definitions associated with \mathbb{N} .

Defining functions between families The prime example of this can be found in the generalization of entailment of propositions to maps of finite sets families.

Definition 5.3.17. Given a pair of families $\sigma_0 : X_0^\bullet \multimap X$ and $\sigma_1 : X_1^\bullet \multimap X$, a map from σ_0 to σ_1 is a map $f : X_0^\bullet \rightarrow X_1^\bullet$ such that $\sigma_1 \circ f = \sigma_0$.

Exercise 5.9. Show that a map of families $f : \sigma_0 \rightarrow \sigma_1$ is equivalent to a dependent map $(x : X) \rightarrow \sigma_0^{-1}(x) \rightarrow \sigma_1^{-1}(x)$.

Example 5.3.18. Consider the set of functions from $\text{id} : X \multimap X$ to $\sigma : X^\bullet \multimap X$ are the set-theoretic counterpart to $\text{Tm}(\Gamma, A)$. Under this analogy, X is the “context” and σ is the “type”. Finally, using Exercise 2.2 we see that $\text{id} \Rightarrow \sigma$ corresponds to terms of type σ .

Unfortunately, there is no straightforward way to lift this definition to operate on equivalence classes $[\sigma_0]$ and $[\sigma_1]$. In other words, there is no straightforward way to define a set $x \Rightarrow y$ of maps between $x, y : \mathbb{N}$. Observe that while we would like to send would send $[\sigma_0]$ and $[\sigma_1]$ to the set of functions $\sigma_0 \rightarrow \sigma_1$, this does not respect equivalence of families. In particular, the set of functions from $\sigma_0 \rightarrow \sigma_1$ is distinct from that of $\sigma'_0 \rightarrow \sigma_1$ even if $\sigma_0 \cong \sigma'_0$; we merely send equivalent families to isomorphic sets.

In order to salvage this definition, we could capitalize on the axiom of choice. After all, the problem is just that we have too many possibilities for $p \Rightarrow q$ and one can always simply revert to making an arbitrary choice. We could make this more palatable by only making a single arbitrary choice: a particular element of the equivalence class $[\pi : \mathbb{N}^\bullet \multimap \mathbb{N}]$. If we pick a specific realization of π rather than just the equivalence class provided by the mapping-in property, it becomes straightforward to define \Rightarrow :

$$n \Rightarrow m = \pi^{-1}(n) \rightarrow \pi^{-1}(m)$$

The root of the issue is that this definition is simply not invariant in the equivalence class of π . As a result, in order to define \Rightarrow we cannot simply rely on the API provided by the mapping in property governing \mathbb{N} . We must instead use a secondary piece of data: a particular choice of representative of the equivalence class $[\pi]$.

This story plays out again once we shift to the more complex universes necessary to model \mathbf{U} . In such cases, the best we can hope for is to represent families of *small sets* $\text{Fam}(X)$ where *small* is some size cut-off larger than “finite” but still smaller than “arbitrarily large”. Once more we find ourselves unable to define important operations on $\text{Fam}(X)$ and, accordingly, without a suitably strong universal characterization of \mathbf{U} . In fact, the story is typically even more fractured for \mathbf{U} . We typically do not even ensure that types have a unique representative (Corollary 5.3.15). However, even if we charitably assume that this property holds, the above discussion shows that it would not suffice for our desired applications.

All of this is in contrast to Ω (or, indeed, every other connective of type theory) where the mapping-in or mapping-out property determined every other operation we wished to define. While we can work around this deficiency by making an arbitrary choice, it begs the question of whether a version of \mathbb{N} can be equipped with an equally powerful mapping-in property allowing us define all the operations we expect on universes.

5.3.4 The groupoid of finite sets

We have two potential maneuvers available if we wish to alter $X \rightarrow \mathbb{N}$: either change the meaning of \rightarrow or change the definition of \mathbb{N} . In fact, we will do both: we will switch from viewing the universe of finite sets as a set to viewing it as a *groupoid*. In so doing, we will also replace/reinterpret \rightarrow to mean “map of groupoids” rather than “map of sets”.

We begin by noticing that the key problem with defining \Rightarrow using only the mapping-in property of \mathbb{N} (i.e., using $\mathbf{Fam}_{\text{fin}}(X)$) was that equivalent families must be treated identically. This would theoretically be fine, except that a pair of families might be equivalent in multiple ways: $\{0, 1\} \rightarrow \{\star\}$ is equivalent to itself via both swap and id.

Accordingly, we wish for an replacement of \mathbb{N} and, with it, $\mathbf{Fam}_{\text{fin}}(X)$ which can account for these multiple different identifications without collapsing all of them into a single equality. This is precisely what a groupoid is for: equipping a set of objects with a proof-relevant equality relation allowing us to state that two objects may be equivalent in multiple distinct ways.

In other words, we will replace \mathbb{N} with a groupoid $\mathbb{N}^{\text{grp d}}$. The resulting mapping-in property will then state that $X \rightarrow \mathbb{N}^{\text{grp d}}$ is equivalent to the *groupoid* $\mathbf{Fam}_{\text{fin}}^{\text{grp d}}(X)$. The groupoid structure of the latter will keep track of the multiple distinct identifications of isomorphic families and thereby allow us to define operations like \Rightarrow without arbitrary choices. We will begin by defining $\mathbb{N}^{\text{grp d}}$.

Definition 5.3.19. $\mathbb{N}^{\text{grp d}}$ is defined such that (1) the carrier set is exactly \mathbb{N} and (2) the set of equivalences $\mathcal{R}_{\mathbb{N}}(n, m)$ is the set of bijections from n element sets to m element sets (and therefore empty if $n \neq m$).

Exercise 5.10. Equip $\mathbb{N}^{\text{grp d}}$ with the remaining structure of a groupoid (\circ , id , $-^{-1}$) and prove that it satisfies the necessary laws.

While the underlying set of $\mathbb{N}^{\text{grp d}}$ is the same, the non-trivial identifications ensure track all the distinct ways of identifying finite sets. For instance, a groupoid homomorphism $f : \mathbf{1} \rightarrow \mathbb{N}^{\text{grp d}}$ which picks out 2 can be identified in two distinct ways: either using the identify isomorphism or using $\text{swap} \in \mathcal{R}_{\mathbb{N}}(2, 2)$.

Our next goal is to characterize the groupoid of homomorphisms $X \rightarrow \mathbb{N}^{\text{grp d}}$. As one might guess from the definition of $\mathbb{N}^{\text{grp d}}$, we shall show that it is equivalent to a groupoid enhancement $\mathbf{Fam}_{\text{fin}}^{\text{grp d}}(X)$ of $\mathbf{Fam}_{\text{fin}}(X)$.

Definition 5.3.20. Given a set X , define $\mathbf{Fam}_{\text{fin}}^{\text{grp d}}(X)$ to be the groupoid whose carrier set is the set finite families over X and such that $\mathcal{R}(\sigma_0, \sigma_1)$ is the set of equivalences of finite families $\sigma_0 \cong \sigma_1$ over X .

Exercise 5.11. Fill out the remaining structure of $\mathbf{Fam}_{\text{fin}}^{\text{grp d}}(X)$ by describing id , \circ , and $-^{-1}$.

We now state the mapping-in property for $\mathbb{N}^{\text{grp d}}$:

Theorem 5.3.21. Given a set X , there is an equivalence of groupoids between groupoid homomorphisms $\Delta X \rightarrow \mathbb{N}^{\text{grp d}}$ and $\mathbf{Fam}_{\text{fin}}^{\text{grp d}}(X)$.

Let us see how the richer $\mathbf{Fam}_{\text{fin}}^{\text{grp d}}(X)$ actually addresses the problem we set out to solve. Remembering Example 4.3.13, we can realize \Rightarrow as a *groupoid homomorphism* $\Rightarrow : \mathbf{Fam}_{\text{fin}}^{\text{grp d}}(X) \times \mathbf{Fam}_{\text{fin}}^{\text{grp d}}(X) \rightarrow S$:

$$\sigma_0 \Rightarrow_0 \sigma_1 = \sigma_0 \rightarrow \sigma_1 \qquad \iota_0 \Rightarrow_1 \iota_1 = \lambda f \rightarrow \iota_1 \circ f \circ \iota_0^{-1}$$

Crucially, we are now merely required to send equivalent families to *isomorphic* sets instead of having to produce identical sets for equivalent families. All told: we recover exactly the story we had for propositions. There is a universal property controlling $\mathbb{N}^{\text{grp d}}$ and it suffices to construct everything we desire on $\mathbb{N}^{\text{grp d}}$ *qua* a universe of finite sets without making any additional choices.

In particular, the universe is now described in a similar manner to other connectives of Section 2.4 with all operations on the universe flowing from a single mapping-in property. The only catch is that in order to achieve this feat, we required a proof-relevant notion of equality and, consequently, the universe of finite sets is neither finite (expected) nor a set (unexpected).

5.3.5 From groupoids to univalence

Let us now note that in search of a good mapping-in property, we have stumbled upon a semantic model of univalence.

Theorem 5.3.22. Fix a pair of families σ_0, σ_1 over X and write $f_0, f_1 : \Delta X \rightarrow \mathbb{N}^{\text{grp d}}$ to the corresponding maps into $\mathbb{N}^{\text{grp d}}$. The following two sets are in bijection:

1. The set of equivalences $\sigma_0 \cong \sigma_1$ over X

2. The set $\mathcal{R}_{\Delta X \rightarrow \mathbb{N}^{\text{grpd}}}(f_0, f_1)$

Proof. This is a direct consequence of Theorem 5.3.21; the equivalence of groupoids $\iota : \mathbf{Fam}_{\text{fin}}^{\text{grpd}}(X) \simeq (\Delta X \rightarrow \mathbb{N}^{\text{grpd}})$ induces a bijection:

$$\mathcal{R}_{\mathbf{Fam}_{\text{fin}}^{\text{grpd}}(X)}(\sigma_0, \sigma_1) \cong \mathcal{R}_{\Delta X \rightarrow \mathbb{N}^{\text{grpd}}}(\iota(\sigma_0), \iota(\sigma_1)) \quad \square$$

In other words, in order to obtain a richer mapping-in property for \mathbb{N} , we were forced to upgrade to \mathbb{N}^{grpd} and we therefore passed from the *isomorphism reflection* result (Corollary 5.3.15) to the actual equivalence required by univalence.

With more effort, we are able to construct a groupoid serving as the universe of small sets analogues to \mathbb{N}^{grpd} and prove a version of Theorem 5.3.22 for this larger universe. Using this result along with the model of intensional type theory in groupoids sketched in Section 4.3, we are able to prove the following result:

Theorem 5.3.23. *Intensional type theory with one univalent universe is consistent.*

We remark that it is not possible to extend the groupoid model to account for nested univalent universes. We can actually prove this to be the case using a pair of previously discussed results: all types in the groupoid model satisfy $U(\text{IP})^2$ (Section 4.3.2.1) and in a hierarchy of nested univalent universes U_i must not satisfy $U(\text{IP})^{i+1}$ (Theorem 5.2.15). Combining these two results, we see that there can be no pair of univalent universes $U_0 : U_1$ within the groupoid model.

In order to find a semantic model with two nested univalent universes, we require a generalization of groupoids allowing for a proof-relevant equality *between proofs of equalities*. That is, we must find a model which refutes $U(\text{IP})^2$. Such a model using 2-groupoids, which generalize groupoids in much the same way that groupoids generalize sets. This pattern can be continued indefinitely: there is a model of intensional type theory in n -groupoids which supports a length n hierarchy of univalent universes. In order to justify the consistency of full HoTT, we therefore must pass to the limit of this process and consider so-called ∞ -groupoids or, as they are sometimes known, homotopy types.

We have seen that searching for a well-adapted mapping-in property for a universe leads to univalence, but what can be said about a universe which is merely known to satisfy the univalence axiom? In fact, it is believed that there is a precise correspondence between models of homotopy type theory and universes satisfying mapping properties akin to \mathbb{N}^{grpd} . This is a consequence of the deeper assertion that homotopy type theory is the internal language of elementary ∞ -topoi. In particular, we should view a good universal property of universes as a moral reformulation of univalence.

All this leads us to the startling realization originating with Hofmann and Streicher [HS98] and culminating in the univalence axiom [Voe10]: accounting for universes properly requires dispensing with UIP and shifting to models based on ∞ -groupoids and homotopy theory [AW09; Voe06].

5.4 Cubical type theory

Thus far in this chapter, we have introduced the univalence axiom and studied a few of its consequences. Hopefully the reader has been convinced that this is an interesting principle with which to extend type theory and that it at least offers partial compensation for the loss of the extensional equality type. However, so far we have considered only the extension of ITT by an simple axiom to obtain univalence and, consequently, the resulting theory does not satisfy canonicity.

In particular, it is not difficult to encounter interesting closed elements of type \mathbf{Nat} which are constructed via univalence, but in core HoTT these programs cannot be evaluated to closed numerals. Famously, Brunerie [Bru18] gave a concise construction of an element of the type $\sum_{n:\mathbf{Nat}} \pi_4(S^3) = \mathbb{Z}/n\mathbb{Z}$ but the lack of canonicity meant that actually working out the concrete $n : \mathbf{Nat}$ for which this equation held was considerably more difficult [Bru16]. This is far from the only example: the proof that $\pi_1(S^1) \simeq \mathbb{Z}$ referenced in Section 5.2 ought to give an algorithm for computing the *winding number* of a map $S^1 \rightarrow S^1$, but this algorithm can only be run if canonicity holds.

Remark 5.4.1. In fact, in Section 5.2 we assumed function extensionality along with univalence. A more careful account would allow us to derive the former from the latter and in fact our solution canonicity and univalence will handle *funext en passant*. \diamond

At first, one might hope that this problem can be fixed “locally”. Perhaps one can simply add a definitional equality to *ua* to recover canonicity. Unfortunately, no such obvious equalities present themselves. A moment’s contemplation will reveal how while there is a reasonable candidate for *transp* applied to *ua*(...), the general case of \mathbf{J} and *ua* is far murkier; such an equation must correctly handle, for instance, the application of *sym* and *trans* to *ua* along with any other number of constructions. More generally, we justified our definition of \mathbf{Id} around the idea that every element of $\mathbf{Id}(A, a, b)$ was controlled by *refl*, but this is simply no longer the case in the presence of *ua*.

Accordingly, our approach to balancing canonicity alongside univalence will involve be a more global and radical change. We shall reimagine the intensional identity type in order to give it a new mapping in property which gives us the flexibility we need to implement univalence. The result of these changes will be cubical type theory [CCHM18; AFH18; Ang+21].

Unfortunately, cubical type theory is vastly more complex than any other type theory we have discussed in this book. Accordingly, we cannot realistically present in the same detail that we have given to ETT or ITT. Our compromise is to introduce what we term *core cubical type theory* in this section. We detail the required modifications to the judgmental structure of type theory, present the additional operations necessary to manipulate them, and sketch how these operations behave and can be used to implement univalence. The last step, however, will mostly be cursory and we will omit most of the rules governing

these operations. We do, however, return to them in Section 5.5 where we discuss some of these details more thoroughly (though still not in the entirety). Our goal is to provide a working knowledge of cubical type theory, rather than a precise account. For the latter, we refer the reader to Angiuli et al. [Ang+21] which does include a more exhaustive account of the theory.

The basis of cubical type theory In this section, we discuss the rules that must be added to intensional type theory in order to arrive at cubical type theory. For concreteness, we will take our base type theory to type theory without any sort of identity type. It is possible to include the intensional identity type as it is possible to extend cubical type theory with indexed inductive types more generally. However, we shall set about to find a better behaved identity type (*path types*) and so its inclusion is superfluous.

5.4.1 *A judgmental structure for identity types*

We begin by convincing ourselves that the judgmental structure of cubical type theory is, in fact, helpful for our problem of giving the identity type a mapping-in property. We begin by observing that we have already attempted to provide an identity type with such a characterization: this was the extensional identity type we moved away from in Chapter 4. There is not an obvious alternative judgmental structure in intensional type theory for the identity type to internalize, so we shall invent one.

This entire process will be broken up into two steps:

1. introduce a new form of judgment and define the new identity type to internalize it,
2. equip each type with additional operations such that this new identity type can implement the expected operations.

We shall eventually see that the first step occupies our attention in Section 5.4.2, while the second takes up Sections 5.4.4 and 5.4.6

It is notable that these two steps are actually distinct: with both the intensional and extensional identity types, once we fixed the judgmental structure we internalized all the rules of the identity type came more-or-less for free. In fact, the same will be true here: the second step does not alter the behavior of the identity type per se. The issue is that the judgmental structure being internalized is no longer definitional equality and so we must add additional structure to all types in order to ensure that this new structure is a useful approximation of equality.

More heuristically, we cannot internalize actual judgmental equality via a mapping-in property and so we internalize a new judgmental structure for *identifications*. We then attempt to paper over the difference between these new judgmental identifications and

actual definitional equality by equipping every single type with additional operations ensuring the former is closer to the latter.

Notation 5.4.2. With an eye towards cubical type theory, we will refer to our new identity type as a *path type* and write $\mathbf{Path}(A, a, b)$ and occasionally refer to identifications as *paths*.

In particular, it is only after both steps are completed that we will have a replacement for \mathbf{Id} that we can contemplate using for univalence. There is a degree of flexibility in how we draw the line between these two steps in cubical type theory. We can make the judgmental structure relatively light-weight for the precise of making the operations on types more onerous or vice versa. This division is the source of the differences between the various flavors of cubical type theory, but overall the differences are slight. We will choose to follow Angiuli et al. [Ang+21] and adopt a relatively minimal judgmental structure at the expense of slightly more complex operations on types.

Let us warm up by considering a direct approach following Licata and Harper [LH12] loosely. We need a new judgment to internalize identifications, so let us simply introduce a new sort of *identifications* α, β, γ which reify identifications and a new judgment $\Gamma \vdash \alpha : a = b : A$ stating that α is such an identification between $a, b : A$. As usual, we will write $\text{Id}(\Gamma, a, b, A)$ for the set of α satisfying $\Gamma \vdash \alpha : a = b : A$. The idea is that we can now at least easily define $\mathbf{Path}(A, a, b)$ via the follow natural isomorphism:

$$\text{Tm}(\Gamma, \mathbf{Path}(A, a, b)) \cong \text{Id}(\Gamma, a, b, A)$$

This completes our goal of defining $\mathbf{Path}(A, a, b)$ and it yields all the necessary rules for this type. The reader will immediately notice, however, that this type is impossible to use and absolutely not a substitute for the identity type. Indeed, just because we claimed that $\Gamma \vdash \alpha : a = b : A$ reifies identifications does nothing to actually force α to behave like any sort of equality. We have only shifted the work into specifying this judgment. For instance, we might choose to include a “reflexivity identification” via the following rule:

$$\frac{\Gamma \vdash a = b : A}{\Gamma \vdash \mathbf{reflId} : a = b : A} \text{📎}$$

Of course, this cannot be the only rule governing our new judgment; the point of this exercise was to allow for additional identifications (such as ua) to arise naturally. In order to do this, we can simply add other inference rules to this judgment! While we do not detail the process here, the reader can imagine that *e.g.*, an identification of pairs can be constructed from identifications of components.

These rules ensure that we can construct elements of $\mathbf{Path}(A, a, b)$, but do not actually give us much leverage in *using* elements of this new type. Our elimination rule for $\mathbf{Path}(A, a, b)$ lets us conclude that there is some identification between a and b , but this is of limited use: there is nothing like \mathbf{J} for identifications or even the equivalent of subst .

Before when identity types internalized definitional equality, we relied on the fact that everything in type theory was automatically congruent and substitutive with respect to definitional equality. Now we are internalizing *identifications* and nothing forces types in our theory to respect identifications in the same way. This is what the second step of the procedure above referred to: we will require additional operations on terms to bridge this gap. For instance, for each type family $\Gamma.A \vdash B$ type there must be an equivalent of subst which sends identifications $\Gamma \vdash \alpha : a = b : A$ to maps between $B(a)$ and $B(b)$.

However, we will not attempt to unfold this further. The problem is that it is difficult to present the full set of rules governing $\Gamma \vdash \alpha : a = b : A$ as well as to present the set of operations all types must enjoy in order to force them to respect identifications. The first problem is the most serious and stems from our desire to support univalence. If we are to have univalence, then we know that there will be elements $a, b : A$ such that the collection of identifications between a and b contains distinct elements and, accordingly, we will quickly run into the need for non-trivial identifications between identifications.

In fact, one can imagine these arising even without univalence: we had discussed that a pair of identifications $\Gamma \vdash \alpha : a = a' : A$ and $\Gamma \vdash \beta : b' = b : B$ ought to induce an identification $\Gamma \vdash (\alpha, \beta) : \mathbf{pair}(a, b) = \mathbf{pair}(a', b') : A \times B$ and we ought to arrange that $(\mathbf{reflId}, \mathbf{reflId}) = \mathbf{reflId}$. To properly account for this and other “higher identifications”, we are quickly led to introducing a new judgment for governing identifications between identifications. As the reader might guess, however, the problem does not stop here and we require judgments for identifications between identifications between identifications... This infinite regress then becomes plain. Accordingly, rather than special-casing a judgment for identifications between terms, we shall design an apparatus which smoothly handles identifications of arbitrary “height”.

It is here that we encounter cubes for the first time. Cubical type theory starts from the insight that an identification between $a, b : A$ can be viewed as a function $\mathbb{I} \rightarrow A$ from some “type” \mathbb{I} . Since we already have a good idea of how to account for functions within the judgments of type theory, if we could recast identity types more into the shape of functions we could reuse this knowledge.

Of course, it is not obvious that functions and identity types share much in common.⁶ A small amount of topological intuition can help motivate this approach: we can say that two points in a space $x, y \in X$ are path-connected just when there is a continuous function from the real interval $p : [0, 1] \rightarrow X$ such that $p(0) = x$ and $p(1) = y$. The geometry of $[0, 1]$ ensures that this notion of identification is actually an equivalence relation. For instance, transitivity comes from the map $[0, 1] \rightarrow [0, 1] \vee [0, 1]$ dividing the interval into two halves and the continuity of $1 - x$ provides symmetry. A major advantage of this definition is that it stacks to identifications between identifications without additional effort: we just take functions from $[0, 1] \times [0, 1]$ satisfying the relevant boundary conditions.

⁶We have already seen hints of this in Section 5.2.3 with the higher-inductive type for the interval

Of course, we have nothing like the real interval in type theory, nor do we intend to add it. However, we can add a judgmental structure which simulates some of its properties and use this as the basis for our definition of an identification in a type. We shall add a faux type \mathbb{I} to our theory and extend our grammar of context to hypothesize over “variables” of type \mathbb{I} such that an identification is then just a term in a context containing such an interval variable.

Remark 5.4.3. It is not yet clear why \mathbb{I} must be a separate structure rather than an ordinary type. Indeed, this is a subtle point and relates to the additional operations necessary to implement the equivalent \mathbf{J} and its related operations. In fact, we shall see that \mathbb{I} cannot support these operations and so it cannot be a type. However, in all other respects it *does* behave like a type: we shall see that the substitution calculus around \mathbb{I} as well as the rules for forming elements its elements are essentially the same as for terms. For this reason, one often refers to \mathbb{I} as a *pre-type*. \diamond

5.4.2 The interval and its structure

Let us make this discussion more formal. We introduce a new judgment $\Gamma \vdash r : \mathbb{I}$ which signifies that r is an element of this interval “pre-type” and has the presupposition $\vdash \Gamma \text{ cx}$. We then introduce a new form a context stating that one may hypothesize over elements of \mathbb{I} . All told, the rules for this are given as follows:

$$\begin{array}{c}
\frac{\vdash \Gamma \text{ cx}}{\vdash \Gamma.\mathbb{I} \text{ cx}} \quad \frac{\vdash \Gamma \text{ cx}}{\Gamma.\mathbb{I} \vdash \mathfrak{p} : \Gamma} \quad \frac{\vdash \Gamma \text{ cx}}{\Gamma.\mathbb{I} \vdash \mathfrak{q} : \mathbb{I}} \quad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash r : \mathbb{I}}{\Delta \vdash r[\gamma] : \mathbb{I}} \\
\\
\frac{\Delta \vdash \gamma : \Gamma \quad \Delta \vdash r : \mathbb{I}}{\Delta \vdash \mathfrak{p} \circ \gamma.r = \gamma : \Gamma} \quad \frac{\Delta \vdash \gamma : \Gamma \quad \Delta \vdash r : \mathbb{I}}{\Delta \vdash \mathfrak{q}[\gamma.r] = r : \mathbb{I}} \quad \frac{\Delta \vdash \gamma : \Gamma.\mathbb{I}}{\Delta \vdash \gamma = (\mathfrak{p} \circ \gamma).\mathfrak{q}[\gamma] : \Gamma.\mathbb{I}} \\
\\
\frac{\Gamma_1 \vdash \gamma_2 : \Gamma_2 \quad \Gamma_0 \vdash \gamma_1 : \Gamma_1 \quad \Gamma_2 \vdash r : \mathbb{I}}{\Gamma_0 \vdash r[\gamma_2 \circ \gamma_1] = r[\gamma_2][\gamma_1] : \mathbb{I}} \quad \frac{\Gamma \vdash r : \mathbb{I}}{\Gamma \vdash r[\mathbf{id}] = r : \mathbb{I}}
\end{array}$$

We shall $\Gamma \vdash r : \mathbb{I}$ as a *dimension term* and \mathfrak{q} as a *dimension variable*.

Notation 5.4.4. We write $\gamma.\mathbb{I}$ for the analogous substitution to $\gamma.A$.

Exercise 5.12. Define $\text{Int}(\Gamma)$ to be the set $\{r \mid \Gamma \vdash r : \mathbb{I}\}$. Rephrase the above substitution rules and equalities using $\text{Int}(\Gamma)$ and, in particular, isolate a mapping-in property for $\Gamma.\mathbb{I}$.

Notation 5.4.5. The reader will notice that while context extension with an interval is formally distinct from $\Gamma.A$, the substitution calculus is the same around both. Consequently,

it is not difficult to adapt the translation from named variables to formal syntax with explicit substitutions to account for interval “variables”. When we write informal programs in cubical type theory, we shall therefore use essentially the same notation for variables of \mathbb{I} as we have for variables of a type A . By convention, we shall use the letters i, j, k for these *dimension variables*.

All told then, \mathbb{I} has been added to our theory such that it behaves more-or-less like a type without any introduction or elimination rules and we can only hypothesize variables of type \mathbb{I} and pass them around. This is far less structure than the real interval $[0, 1]$, but it is already almost enough to realize our judgmental structure for identifications: an identification in A can simply be taken as an element $\Gamma.\mathbb{I} \vdash a : A[p]$. What we are missing is some means of stating what, precisely, is being identified by such an a . In the topological case, an identification was a continuous function $p : [0, 1] \rightarrow X$ which identified $p(0)$ with $p(1)$. Accordingly, we now augment \mathbb{I} with two closed dimension terms $0, 1$ and understand $\Gamma.\mathbb{I} \vdash a : A[p]$ to be identifying $a[\mathbf{id}.0]$ and $a[\mathbf{id}.1]$.

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash 0, 1 : \mathbb{I}}$$

Lemma 5.4.6. *For every $\Gamma \vdash a : A$ there is an identification of a with itself.*

Proof. Given such an a , the term $p = a[p]$ is an element of $\Gamma.\mathbb{I} \vdash A[p]$ type and it is routine to check that $p[\mathbf{id}.0] = a = p[\mathbf{id}.1]$. \square

We have used \mathbb{I} to recover the bespoke identification judgment from before and in a less ad-hoc manner. Just as before, we may define a path type to internalize this new structure directly:

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma \vdash a, b : A}{\Gamma \vdash \mathbf{Path}(A, a, b) \text{ type}}$$

$$\frac{\Gamma \vdash a, b : A \quad \Gamma.\mathbb{I} \vdash p : A[p] \quad \Gamma \vdash a = p[\mathbf{id}.0] : A \quad \Gamma \vdash a = p[\mathbf{id}.1] : A}{\Gamma \vdash \lambda(p) : \mathbf{Path}(A, a, b)}$$

$$\frac{\Gamma \vdash p : \mathbf{Path}(A, a, b) \quad \Gamma \vdash r : \mathbb{I}}{\Gamma \vdash \mathbf{papp}(p, r) : A} \quad \frac{\Gamma \vdash p : \mathbf{Path}(A, a, b)}{\Gamma \vdash \mathbf{papp}(p, 0) = a : A \quad \Gamma \vdash \mathbf{papp}(p, 1) = b : A}$$

$$\begin{array}{c}
\Gamma \vdash a, b : A \\
\hline
\Gamma.\mathbb{I} \vdash p : A[p] \quad \Gamma \vdash a = p[\mathbf{id}.0] : A \quad \Gamma \vdash a = p[\mathbf{id}.1] : A \quad \Gamma \vdash r : \mathbb{I} \\
\hline
\Gamma \vdash \mathbf{papp}(\lambda(p), r) = p[\mathbf{id}.r] : A \\
\hline
\Gamma \vdash p : \mathbf{Path}(A, a, b) \\
\hline
\Gamma \vdash \lambda(\mathbf{papp}(p[p], q)) = p : \mathbf{Path}(A, a, b)
\end{array}$$

Notice that, unlike ordinary function types, path types specify their behavior on certain arguments in their types. In particular, path types specify the “start” and “end” elements (the two terms being identified). The rules for path types then record this by requiring that elements of $\mathbf{Path}(A, a, b)$ are not just arbitrary elements of A depending on \mathbb{I} but that they additionally satisfy certain *boundary conditions* at 0 and 1.

Exercise 5.13. Show that the above rules are precisely equivalent to requiring the following mapping-in property for $\mathbf{Path}(A, a, b)$:

$$\mathrm{Tm}(\Gamma, \mathbf{Path}(A, a, b)) \cong \{p \in \mathrm{Tm}(\Gamma.\mathbb{I}, A[p]) \mid p[\mathbf{id}.0] = a \wedge p[\mathbf{id}.1] = b\}$$

Exercise 5.14. Use Lemma 5.4.6 to define an element $\mathbf{refl}(a) : \mathbf{Path}(A, a, a)$ for every $\Gamma \vdash a : A$.

Notation 5.4.7. For expository purposes, it is also helpful to have a type $\mathbf{\Pi}(\mathbb{I}, A)$ with the following mapping-in property:

$$\mathrm{Tm}(\Gamma, \mathbf{\Pi}(\mathbb{I}, A)) \cong \mathrm{Tm}(\Gamma.\mathbb{I}, A)$$

We shall not regard this as part of our official definition of cubical type theory and use it only for small informal examples. In these few occurrences of this “ $\mathbf{\Pi}$ -type”, we shall use the ordinary syntax for functions, using the observation above that we can translate “named interval variables” into the formal substitution calculus for \mathbb{I} .

5.4.3 Cultivating intuition for path types

Before proceeding to the other rules of cubical type theory, we take a moment to explore the consequences of including the interval within type theory. For this, and in cubical type theory more generally, it is helpful to use a small amount of topological intuition, guided by the observation that a term $\mathbf{1}.\mathbb{I} \dots \mathbb{I} \vdash a : A[p^n]$ which depends on n copies of \mathbb{I} can be visualized as an n -dimensional cube in A . In low dimension, we therefore have points in A , lines in A , squares in A , and cubes in A for $n = 0, 1, 2, 3$ respectively. Let us illustrate the

case where $n = 2$ more thoroughly. Given $\mathbf{1}.\mathbb{I} \vdash a : A[p^2]$, the two dimension variables serve as “axes” for this square and so we can “draw” a as the following square:

$$\begin{array}{ccc}
 a[\mathbf{id}.0.0] & \xrightarrow{a[\mathbf{id}.0]} & a[\mathbf{id}.1.0] \\
 \downarrow a[\mathbf{id}.0.\mathbb{I}] & & \downarrow a[\mathbf{id}.1.\mathbb{I}] \\
 a[\mathbf{id}.0.1] & \xrightarrow{a[\mathbf{id}.1]} & a[\mathbf{id}.1.1]
 \end{array}$$

The four closed terms one obtains by specializing a with the four substitutions $\mathbf{1} \vdash \mathbf{id}.\epsilon.\epsilon' : \mathbf{1}.\mathbb{I}$ are the vertices. Next, there are four substitutions from $\mathbf{1}$ to $\mathbf{1}.\mathbb{I}$ which implement the first or second \mathbb{I} with a constant and the other \mathbb{I} with \mathfrak{q} . Applying each of these substitutions to a yields the edges of the square. Finally, a itself is the entire square.

We have chosen to draw this square with the leftmost \mathbb{I} in $\mathbf{1}.\mathbb{I}$ as the horizontal axis and the second as the vertical axis and such that the horizontal axis grows from to the right while the vertical axis grows down. This convention is reasonably standard—it matches the typical orientation of commutative diagrams in category theory—but it is often helpful to disambiguate these diagrams by using named variables and labeling axes. For instance, we might have written $i : \mathbb{I}, j : \mathbb{I} \vdash a(i, j) : A$ and depicted the above square as follows:

$$\begin{array}{ccc}
 \begin{array}{c} i \\ \rightarrow \\ j \downarrow \end{array} & \begin{array}{ccc} a(0,0) & \xrightarrow{a(i,0)} & a(1,0) \\ \downarrow a(0,j) & & \downarrow a(1,j) \\ a(0,1) & \xrightarrow{a(i,1)} & a(1,1) \end{array}
 \end{array}$$

Remark 5.4.8. We note that in the above example we have assumed that A does not depend on either dimension variable but this restriction is not mandatory. We will have occasion to study such *heterogeneous* squares at various points. \diamond

This schematic visualization highlights one of the major benefits of using \mathbb{I} to structure identifications compared to a direct judgment $\Gamma \vdash \alpha : a = b : A$: we can now seamlessly account for identifications between identifications simply by adding more than one copy of \mathbb{I} to the context. Moreover, path types between path types of A are really no more complex to manipulate than ordinary path types as both are simply kinds of functions valued in A .

There is another major benefit to using \mathbb{I} : we have no need to add further rules of \mathbb{I} to customize the behavior of path types in each connective. For instance, there is no need

for a rule that “identifications in a pair can be built from a pair of identifications”. This fact is already derivable from those rules governing dependent sums generally. In fact, path types enjoy a number of remarkable extensionality principles (including function extensionality) without additional effort on our part.

This trace back to a subtle point: when we isolated identifications as a new judgment, nothing connected it to the behavior of types or terms. Here, however, we have smuggled identifications in through the existing apparatus of contexts and substitutions and so the existing equations for types and terms automatically apply to identifications.

For instance, the η law for dependent sums states that $\text{Tm}(\Gamma, \Sigma(A, B))$ is isomorphic to $\sum_{a \in \text{Tm}(\Gamma, A)} \text{Tm}(\Gamma, B[\text{id}.a])$. If we choose $\Gamma = \Gamma_0.\mathbb{I}$ and specialize to the case where $B = B_0[\mathbf{p}]$ for simplicity, this immediately yields the following:

Lemma 5.4.9. *There is a natural bijection of the following shape:*

$$\begin{aligned} \text{Tm}(\Gamma_0, \text{Path}(\Sigma(A, B_0[\mathbf{p}]), x, y)) \\ \cong \text{Tm}(\Gamma_0, \text{Path}(A, \text{fst}(x), \text{fst}(y))) \times \text{Tm}(\Gamma_0, \text{Path}(A, \text{snd}(x), \text{snd}(y))) \end{aligned}$$

Exercise 5.15. Prove Lemma 5.4.9

Note that while we have specialized to the simpler case of non-dependent Σ -types, it is only for notational convenience. Even more striking is the case for dependent products.

Lemma 5.4.10. *There is a natural bijection of the following shape:*

$$\text{Tm}(\Gamma, \text{Path}(\Pi(A, B), f, g)) \cong \text{Tm}(\Gamma.A, \text{Path}(B, \text{app}(f[\mathbf{p}], \mathbf{q}), \text{app}(g[\mathbf{p}], \mathbf{q})))$$

In other words, function extensionality is automatically true for path types.

Proof. Let us begin by observing that, by the mapping-in property of path types, we can rephrase our goal as the following:

$$\{p \in \text{Tm}(\Gamma.\mathbb{I}, \Pi(A, B)[\mathbf{p}]) \mid \dots\} \cong \{\text{Tm}(\Gamma.A.\mathbb{I}, B[\mathbf{p}]) \mid \dots\}$$

However, we can further apply the mapping-in property for Π -types to replace the left-hand set with $\{p \in \text{Tm}(\Gamma.\mathbb{I}.A[\mathbf{p}], B[\mathbf{p}.A]) \mid \dots\}$. The conclusion follows immediately from the isomorphism of contexts $\Gamma.\mathbb{I}.A[\mathbf{p}] \cong \Gamma.A.\mathbb{I}$ (Exercise 5.16). \square

Exercise 5.16. Prove that if $\Gamma \vdash A$ type then there are mutually inverse substitutions $\Gamma.\mathbb{I}.A[\mathbb{p}] \vdash \tau_0 : \Gamma.A.\mathbb{I}$ and $\Gamma.A.\mathbb{I} \vdash \tau_1 : \Gamma.\mathbb{I}.A[\mathbb{p}]$.

This has certainly improved on our earlier attempt which simply added a new explicit judgment of identifications but the story is cannot stop here. In particular, we still have done nothing to address the link between $\mathbf{Path}(A, a, b)$ and the actual ability to substitute a for b in a type. That is, we have no operation like that of \mathbf{subst} or, more generally, \mathbf{J} . As mentioned earlier, these operations do not come directly from the interval or judgments upon it. Instead, we shall add them more-or-less as constants to our theory and then, to preserve canonicity, add type-specific equations telling us how they compute.

5.4.4 Coercing along paths

We now introduce the first and most fundamental operation of the two operations we shall add to cubical type theory: \mathbf{coe}_A (short for *coerce*). Roughly, this operation ensures that, from the perspective of a type, all elements of the interval are interchangeable and we shall see momentarily that this is precisely what is required to implement a version of \mathbf{subst} for $\mathbf{Path}(A, a, b)$.

The addition of \mathbf{coe}_A also means a change in the status of \mathbb{I} in our type theory. While we have not added any sort of elimination principle for \mathbb{I} , the reader may have noticed that up till this point there was really nothing which distinguished it from \mathbf{Bool} ; the rules we required of \mathbb{I} were a strict subset of those for \mathbf{Bool} . The coercion operation firmly rules out the possibility that $\mathbb{I} = \mathbf{Bool}$: a type depending on \mathbf{Bool} can be quite different over \mathbf{true} and \mathbf{false} which is precisely the possibility excluded by \mathbf{coe} .

Specifically, if $\Gamma.\mathbb{I} \vdash A$ type then $A[\mathbf{id}.r]$ and $A[\mathbf{id}.s]$ are equivalent for every $\Gamma \vdash r, s : \mathbb{I}$. The typing rule for this constant is given as follows:

$$\frac{\Gamma.\mathbb{I} \vdash A \text{ type} \quad \Gamma \vdash r, s : \mathbb{I} \quad \Gamma \vdash a : A[\mathbf{id}.r]}{\Gamma \vdash \mathbf{coe}_A^{r \rightarrow s}(a) : A[\mathbf{id}.s]}$$

$$\frac{\Gamma.\mathbb{I} \vdash A \text{ type} \quad \Gamma \vdash r : \mathbb{I} \quad \Gamma \vdash a : A[\mathbf{id}.r]}{\Gamma \vdash \mathbf{coe}_A^{r \rightarrow r}(a) = a : A[\mathbf{id}.r]}$$

A priori, \mathbf{coe} may seem as though it does little to advance our goal of implementing \mathbf{subst} for $\mathbf{Path}(A, a, b)$. However, suppose we are given a path $\Gamma \vdash p : \mathbf{Path}(A, a, b)$ along with a type $\Gamma.A \vdash C$ type, applying the ordinary rule for substitution, we obtain $\Gamma.\mathbb{I} \vdash C' = C[\mathbb{p}.\mathbf{papp}(p, \mathbb{q})]$ type. Inspection reveals that instantiating C' at 0 and 1 yields $C[\mathbf{id}.a]$ and $C[\mathbf{id}.b]$ and so \mathbf{coe} yields the following operation:

$$\Gamma \vdash \lambda(\mathbf{coe}_{C'[\mathbb{p}.\mathbb{I}]}^{0 \rightarrow 1}(\mathbb{q})) : C[\mathbf{id}.a] \rightarrow C[\mathbf{id}.b]$$

In other words, `coe` can be used to define `subst`. The advantage to `coe` over `subst` is that we can now set about equipping `coe` with a collection of definitional equalities in order to recover canonicity. Unlike `subst`, there shall be no single rule for how `coe` computes in general but, instead, `coeA` will compute depending on the form of A . For example, for closed types such as `Nat`, `U`, or `Bool`, we constrain `coe` with the following:

$$\frac{\Gamma \vdash r, s : \mathbb{I} \quad \Gamma \vdash b : \mathbf{Bool}}{\Gamma \vdash \mathbf{coe}_{\mathbf{Bool}}^{r \rightarrow s}(b) = b : \mathbf{Bool}}$$

Of course, this strategy only works in the simplest example: when the type constructor is closed and cannot depend on the interval in any meaningful way. Most commonly, when A is a type former *e.g.* $\Sigma(B_0, B_1)$, `coeA` will be defined in terms of `coeBi`. In the case of non-dependent case $A = B_0 \times B_1$, for instance, one must add a definitional equality stating $\mathbf{coe}_A^{r \rightarrow s}(a) = \mathbf{pair}(\mathbf{coe}_{B_0}^{r \rightarrow s}(\mathbf{fst}(a)), \mathbf{coe}_{B_1}^{r \rightarrow s}(\mathbf{snd}(a)))$.⁷ In Section 5.5, we shall see that while it is unfeasible to see how univalence ought to compute relative to `J`, it is possible (if difficult) to describe its computation with respect to `coe`.

Our strategy of defining `coeA` in terms of the constituents of A is responsible for another surprising feature of `coe`: if `subst` is defined by instantiating $r = 0$ and $s = 1$, why do we bother to allow for arbitrary r, s ? We shall see that in various situations we require this additional flexibility in order to build up `coe` at more complex types from simpler ones.

We will not detail the equations governing `coe` here, but do provide examples in Section 5.5. Instead, we focus on the equation which leads to the next structure necessary for core cubical type theory: `coe` in `Path(A, a, b)`. At present, we lack the operations necessary to provide an equation specifying how $\mathbf{coe}_{\mathbf{Path}(A, a, b)}^{r \rightarrow s}(p)$ must compute. It is worth sketching the problem informally, so as to properly situate the solution. We wish to formulate a rule of the following shape:

$$\frac{\Gamma.\mathbb{I} \vdash A \text{ type} \quad \Gamma.\mathbb{I} \vdash a, b : A \quad \Gamma \vdash r, s : \mathbb{I} \quad \Gamma \vdash p : \mathbf{Path}(A, a, b)[\mathbf{id}.r]}{\Gamma \vdash \mathbf{coe}_{\mathbf{Path}(A, a, b)}^{r \rightarrow s}(p) = \text{?} : \mathbf{Path}(A, a, b)[\mathbf{id}.s]} \quad \text{🔪}$$

This hole must be filled by a path in A built from `coeA`. The straightforward approach is roughly to “compose” p (the function from \mathbb{I} to $A[\mathbf{id}.r]$) with `coeA r s` (a function $A[\mathbf{id}.r] \rightarrow A[\mathbf{id}.s]$). However, the resulting term does not satisfy the necessary boundary conditions to be an element of `Path(A, a, b)[id.s]`. Instead, we obtain an element of the following:

$$\mathbf{Path}(A[\mathbf{id}.s], \mathbf{coe}_A^{r \rightarrow s}(a), \mathbf{coe}_A^{r \rightarrow s}(b))$$

In other words, we are confronted by the fact that while there is a “line” interpolating between *e.g.*, $\mathbf{coe}_A^{r \rightarrow s}(a)$ and $a[\mathbf{id}.s]$, they are not equal. This mismatch is solved by the

⁷This is often expressed by stating that `coe` is defined “by induction” on the type, but this is misleading. After all, types do not come equipped with any sort of induction principle in general!

second operation for manipulating terms depending on \mathbb{I} : homogeneous composition or **hcomp**. To a first approximation, this operation allows us to take our collection of three lines and compose them into a single path.

However, while the motivating example given above comes from stitching together three sides of a square into a single line, our need to provide type-specific equations for computing this operation in each type forces us to provide a more general composition operator. In order to properly formulate **hcomp** in Section 5.4.6, we begin by extending the judgmental apparatus with the necessary tools to support it.

5.4.5 Cofibrations and faces

Let us fix $1.\mathbb{I} \vdash a : A[\mathbb{p}^2]$ and recall the visualization of a as a square:

$$\begin{array}{ccc}
 a[\mathbf{id}.0.0] & \xrightarrow{a[\mathbf{id}.0]} & a[\mathbf{id}.1.0] \\
 \downarrow a[\mathbf{id}.0.\mathbb{I}] & & \downarrow a[\mathbf{id}.1.\mathbb{I}] \\
 a[\mathbf{id}.0.1] & \xrightarrow{a[\mathbf{id}.1]} & a[\mathbf{id}.1.1]
 \end{array}$$

The edges and vertices in the above square are called the *faces* of a . More generally, a face of a term p is the result from specializing interval variables p depends upon.

The **hcomp** operation which we use to compose paths does so by solving a more general problem. It provides a uniform way to assemble certain collections of matching faces into an entire n -cube. For instance, our earlier desire to combine three lines into a single line can be rephrased into taking three terms representing three faces of a square and extending them to a term representing the entire square.

In general, we should not expect that every matching collection of faces assembles into a cube. For instance, the question of whether a and b are identifiable amounts to asking if a and b are the 0 and 1 faces of some term p . Since we do not expect (or want!) all terms to be identifiable, clearly some subsets of cubes should not always be extendable.

Heuristically, we should be allowed to extend subcubes which are suitably “connected”, but this becomes subtle in higher dimensions. As isolating these well-behaved subcubes is complex, it is helpful to have an judgmental apparatus for isolating particular faces of a given term or type. We do this by introducing a special grammar of propositions which we call *cofibrations*. Informally, these are propositions built from (1) comparing dimension terms for equality and (2) conjunction, disjunction, and universal quantification of \mathbb{I} . We realize this with a new judgment $\Gamma \vdash \alpha \text{ cof}$:

$$\begin{array}{c}
\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \top, \perp \text{ cof}} \quad \frac{\Gamma \vdash \phi, \psi \text{ cof}}{\Gamma \vdash \phi \wedge \psi, \phi \vee \psi \text{ cof}} \quad \frac{\Gamma \vdash r, s : \mathbb{I}}{\Gamma \vdash r = s \text{ cof}} \quad \frac{\Gamma.\mathbb{I} \vdash \phi \text{ cof}}{\Gamma \vdash \forall \phi \text{ cof}} \\
\\
\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash \phi \text{ cof}}{\Delta \vdash \phi[\gamma] \text{ cof}} \quad \frac{\Gamma_1 \vdash \gamma_2 : \Gamma_2 \quad \Gamma_0 \vdash \gamma_1 : \Gamma_1 \quad \Gamma_2 \vdash \phi \text{ cof}}{\Gamma_0 \vdash \phi[\gamma_2 \circ \gamma_1] = \phi[\gamma_2][\gamma_1] \text{ cof}} \\
\\
\frac{\Gamma \vdash \phi \text{ cof}}{\Gamma \vdash \phi[\text{id}] = \phi \text{ cof}}
\end{array}$$

We have omitted the long but unsurprising list of rules shaping how substitutions $\phi[\gamma]$ interact with the various cofibration formers.

In keeping with their obvious relationship propositions, we add another judgment $\Gamma \vdash \phi \text{ true}$ which states that some cofibration ϕ holds in context Γ . For instance, we shall eventually add rules to ensure that $\Gamma \vdash \top \text{ true}$ and $\Gamma \vdash 0 = 0 \text{ true}$ both hold. These rules are, however, complicated by the fact that we also allow one to *hypothesize* the truth of a cofibration in Γ just as one can hypothesize an element of A . Explicitly, given cofibration $\Gamma \vdash \phi \text{ cof}$, we also require a context $\Gamma.\phi$ governed by the following rules:

$$\begin{array}{c}
\frac{\Gamma \vdash \phi \text{ cof}}{\vdash \Gamma.\phi \text{ cx}} \quad \frac{\Gamma \vdash \phi \text{ cof}}{\Gamma.\phi \vdash p : \Gamma} \quad \frac{\Gamma \vdash \Gamma \text{ cof}}{\Gamma.\phi \vdash \phi[p] \text{ true}} \\
\\
\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash \phi \text{ cof} \quad \Gamma \vdash \phi[\gamma] \text{ true}}{\Delta \vdash \phi.\star : \Gamma.\phi} \quad \frac{\Gamma \vdash \phi \text{ cof} \quad \Delta \vdash \gamma : \Gamma.\phi}{\Delta \vdash (\phi \circ p).\star = \gamma : \Gamma.\phi} \\
\\
\frac{\Gamma \vdash \phi \text{ cof}}{\Gamma.\phi \vdash p.\star = \text{id} : \Gamma.\phi}
\end{array}$$

It is helpful to understand $\Gamma.\phi$ as an analog of $\Gamma.A$ but where ϕ is an exceptionally strict form of proposition rather than a full type. For instance, the substitution extension rule for cofibrations $\gamma.\star$ does not allow the user to supply alternative “proofs” or “terms” witnessing that ϕ is true. Instead, it simply requires that the judgment $\Gamma \vdash \phi \text{ true}$ holds and uses \star . In fact, the user is not responsible for providing any evidence whatsoever in their term that $\Gamma \vdash \phi \text{ true}$ holds. In this way, the rule is reminiscent of the conversion rule stating that definitionally equal terms may be exchanged without any explicit instruction by the user: true cofibrations may be inhabited without the user having to provide any explicit witness to their truth.

For this reason, it is apparent that we must maintain strict control over the judgment $\Gamma \vdash \phi \text{ true}$. If this judgment becomes too complex and, for instance, becomes sensitive to what *types* are inhabited in a given context Γ , then it will surely become impossible for

our system to enjoy decidable type-checking. Fortunately, the grammar of cofibrations is sufficiently simple that $\Gamma \vdash \phi$ true is, in fact, decidable. We present a selection of the rules for this judgment below, but we have omitted several. In particular, since $\phi \wedge \psi$ and $\forall \phi$ will not appear in the subset of rules we present for cubical type theory, we will omit their rules here. The reader may trust that they are unsurprising versions of the ordinary rules for propositional logic.

$$\begin{array}{c}
\frac{\Gamma \vdash r = s : \mathbb{I}}{\Gamma \vdash r = s \text{ true}} \quad \frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \top \text{ true}} \quad \frac{\Gamma \vdash \phi, \psi \text{ cof} \quad \Gamma \vdash \phi \text{ true}}{\Gamma \vdash \phi \vee \psi, \psi \vee \phi \text{ true}} \quad \frac{\Gamma \vdash 0 = 1 \text{ true}}{\Gamma \vdash \perp \text{ true}} \\
\\
\frac{\Gamma \vdash \phi \text{ true} \quad \Delta \vdash \gamma : \Gamma}{\Delta \vdash \phi[\gamma] \text{ true}} \quad \frac{\Gamma \vdash \perp \text{ true} \quad \Gamma \vdash \phi \text{ cof}}{\Gamma \vdash \phi \text{ true}} \\
\\
\frac{\Gamma \vdash \phi \vee \psi \text{ true} \quad \Gamma \vdash \xi \text{ cof} \quad \Gamma.\phi \vdash \xi[\mathbb{p}] \text{ true} \quad \Gamma.\psi \vdash \xi[\mathbb{p}] \text{ true}}{\Gamma \vdash \xi \text{ true}} \\
\\
\frac{\Gamma \vdash r, s : \mathbb{I} \quad \Gamma \vdash r = s \text{ true}}{\Gamma \vdash r = s : \mathbb{I}} \quad \frac{\Gamma \vdash \phi, \psi \text{ cof} \quad \Gamma.\phi \vdash \psi[\mathbb{p}] \text{ true} \quad \Gamma.\psi \vdash \phi[\mathbb{p}] \text{ true}}{\Gamma \vdash \phi = \psi \text{ cof}}
\end{array}$$

The three lines of rules are a largely unsurprising recitation of the rules of propositional logic, so we focus on the final line. This final line contains two rules: one reminiscent of equality reflection from Chapter 2 and one akin to very strong propositional univalence principle for cofibrations. That is, the first rule guarantees that if the proposition $r = s$ holds then this can be ‘reflected’ to obtain a definitional equality between r and s . The second rule states that cofibrations which are inter-provable are *definitionally* equal such that, e.g., one may silently exchange $\phi \vee \psi$ and $\psi \vee \phi$ in any term or type.

These last two rules imply that the truth of a cofibration can impact whether or not a term or type is well-formed by, for instance, controlling whether two dimension terms are equal. However, we will also add two principles which much more directly allow cofibrations to influence terms, types, and substitutions. Namely, if $\Gamma \vdash \phi \vee \psi$ true, we will add a rule stating that to e.g., construct a type in Γ it suffices to give a type A_ϕ under the assumption of ϕ and one a second A_ψ under the assumption of ψ such that $A_\phi = A_\psi$ when $\phi \wedge \psi$ is assumed. We require similar rules for terms and substitutions and as well as a twin principle for \perp which simply states that all these judgments collapse if $\Gamma \vdash \perp$ true. These rules are designed to ensure that $\Gamma.\phi \vee \psi$ behaves like the “union” of the contexts $\Gamma.\phi$ and $\Gamma.\psi$. For reasons of space, we give the rules carefully for only types and sketch those for terms:

$$\begin{array}{c}
\frac{\Gamma \vdash \phi, \psi \text{ cof} \quad \Gamma \vdash \phi \vee \psi \text{ true} \quad \Gamma \cdot \psi \vdash A_\psi \text{ type} \quad \Gamma \cdot \psi \wedge \psi \vdash A_\phi[\mathbb{p} \cdot \star] = A_\psi[\mathbb{p} \cdot \star] \text{ type}}{\Gamma \vdash [\phi \hookrightarrow A_\phi \mid \psi \hookrightarrow A_\psi] \text{ type}} \\
\\
\frac{\Gamma \vdash \phi_1, \psi_2 \text{ cof} \quad \Gamma \cdot \phi_1 \vdash A_{\phi_1} \text{ type} \quad \Gamma \cdot \phi_1 \wedge \phi_2 \vdash A_{\phi_1}[\mathbb{p} \cdot \star] = A_{\phi_2}[\mathbb{p} \cdot \star] \text{ type} \quad \Gamma \vdash \phi_i \text{ true}}{\Gamma \vdash [\phi_1 \hookrightarrow A_{\phi_1} \mid \phi_2 \hookrightarrow A_{\phi_2}] = A_{\phi_i}[\mathbf{id} \cdot \star] \text{ type}} \\
\\
\frac{\Gamma \vdash \phi, \psi \text{ cof} \quad \Gamma \vdash \phi \vee \psi \text{ true} \quad \Gamma \vdash A \text{ type}}{\Gamma \vdash [\phi \hookrightarrow A[\mathbb{p}] \mid \psi \hookrightarrow A[\mathbb{p}]] = A \text{ type}} \\
\\
\frac{\Gamma \vdash \perp \text{ true}}{\Gamma \vdash \mathbf{Abort} \text{ type}} \qquad \frac{\Gamma \vdash \perp \text{ true} \quad \Gamma \vdash A \text{ type}}{\Gamma \vdash \mathbf{Abort} = A \text{ type}} \\
\\
\frac{\Gamma \vdash \phi, \psi \text{ cof} \quad \Gamma \vdash \phi \vee \psi \text{ true} \quad \Gamma \vdash A \text{ type} \quad \Gamma \cdot \phi \vdash a_\phi : A[\mathbb{p}] \quad \Gamma \cdot \psi \vdash a_\psi : A[\mathbb{p}] \quad \Gamma \cdot \psi \wedge \psi \vdash a_\phi[\mathbb{p} \cdot \star] = a_\psi[\mathbb{p} \cdot \star] : A[\mathbb{p} \cdot \star]}{\Gamma \vdash [\phi \hookrightarrow a_\phi \mid \psi \hookrightarrow a_\psi] : A} \\
\\
\frac{\Gamma \vdash \perp \text{ true} \quad \Gamma \vdash A \text{ type}}{\Gamma \vdash \mathbf{abort} : A} \qquad \frac{\Gamma \vdash \perp \text{ true} \quad \Gamma \vdash a : A}{\Gamma \vdash \mathbf{abort} = a : A}
\end{array}$$

Advanced Remark 5.4.11. More concisely, these conditions ensure that $\Gamma \cdot \phi \vee \psi$ is a pushout of $\Gamma \cdot \phi$ and $\Gamma \cdot \psi$ over $\Gamma \cdot \phi \wedge \psi$ and that the presheaves for terms, types, *etc.* carry these pushouts to pullbacks. Similarly, they guarantee that $\Gamma \cdot \perp$ is initial and that all relevant presheaves carry this initial object to a terminal object. \diamond

From cofibrations to subcubes These rules finally allow us to deliver on an earlier promise: we can now use cofibrations to isolate particular combinations of faces from a term. Let us consider the context consisting of two dimension variables extended by a cofibration stating either the first variable is 0 or the second is 1:

$$\Gamma = \mathbf{1} \cdot \mathbb{I} \cdot \mathbb{I} \cdot (\mathfrak{q} = 1 \vee \mathfrak{q}[\mathbb{p}] = 0)$$

We know by the above rules for disjunction that giving a term $\Gamma \vdash a : A[\mathbb{p}^3]$ is equivalent to giving a terms $\mathbf{1} \cdot \mathbb{I} \cdot \mathbb{I} \cdot (\mathfrak{q} = 0) \vdash a_0 : A[\mathbb{p}^3]$ and $\mathbf{1} \cdot \mathbb{I} \cdot \mathbb{I} \cdot (\mathfrak{q} = 1) \vdash a_1 : A[\mathbb{p}^3]$ which agree on the overlap. Next, one may use the equality reflection rule for $\mathfrak{q} = 1$ to show that *e.g.*, the substitution $\mathbf{1} \cdot \mathbb{I} \cdot \mathbb{I} \cdot (\mathfrak{q} = 0) \vdash \mathbb{p} \cdot \mathbb{I} \circ \mathbb{p} : \mathbf{1} \cdot \mathbb{I}$ is invertible. We may therefore visualize a_0 and a_1

as lines in A which share a common boundary:

$$\begin{array}{ccc}
 a[\mathbf{id}.0.0] & & \\
 \downarrow a_1 & & \\
 a[\mathbf{id}.0.1] & \xrightarrow{a_0} & a[\mathbf{id}.1.1]
 \end{array}$$

More generally, if ϕ is any cofibration then $\Gamma.\phi \vdash a_\phi : A[\mathbb{p}]$ will consist of some *coherent* collection of faces in A . In other words, ϕ isolates some subset of the faces of an n -cube, and the rule for splitting on disjunctions of cofibrations ensures that a_ϕ consists of a term for each face *such that these terms agree on all overlaps*. The question of whether these faces can be stitched together into a single n -cube amounts to asking whether or not there exists some $\Gamma \vdash a : A$ such that $\Gamma.\phi \vdash a[\mathbb{p}] = a_\phi : A[\mathbb{p}]$. This rephrasing in terms of cofibrations offers two important advantages. First, this formulation has better behavior with respect to substitution: it is clear that any extension in the above sense is stable under substitution and it also ensures that we can sensibly discuss applying substitutions to collections of faces. Second, cofibrations allow us to discuss more exotic faces like the line carved out by the cofibration $i = j$ for two dimension variables i, j . This corresponds to the *diagonal* of a square, rather than any of its standard edges.

Notation 5.4.12. In Section 5.5, we will wish to manipulate cofibrations when working informally with type theory. In general, like dimension variables the substitution calculus ensures that we can largely pretend ϕ is a “type”, but the exceptionally strict properties around cofibrations ensure that we need never actually pass one around. When working informally, we shall therefore treat them in much the same way proof assistants handle implicit arguments: abstracting over them with a bespoke function type (the *partial element type*) but never needing to actually provide explicit terms to apply these function types. We present only the mapping-in property for this type and leave it to the reader to see how ordinary implicit function syntax may be translated to this isomorphism:

$$\mathrm{Tm}(\Gamma, \phi \rightarrow A) \cong \mathrm{Tm}(\Gamma.\phi, A)$$

In the above, $\Gamma \vdash \phi \rightarrow A$ type just when $\Gamma.\phi \vdash A$ type. However, since we shall only use this connective for informal explanations, we will not regard it as part of our definition of cubical type theory and content ourselves with this sketch of its rules.

5.4.6 Composing and filling paths

We are now describe the second operation for manipulating paths **hcomp** and the final component of core cubical type theory. Recall that this operation is intended to take

collections of faces—a subset of an n -cube in A —and assemble them into single n -cube in A . As noted earlier, it is unsound to provide such an operation for *arbitrary* subcubes, but with the apparatus of cofibrations to hand, it is possible to describe a flexible class of shapes for which it is sound: given a term $\Gamma \vdash a_0 : A$ representing an n -cube in A along with a cofibration $\Gamma \vdash \phi \text{ cof}$ and a “ ϕ -partial line” $\Gamma.\phi.\mathbb{I} \vdash a_\phi : A$, which matches a_0 match appropriately, we may glue and extend them using **hcomp** to an $(n + 1)$ -cube in A . The formal rules are as follows with a_0 and a_ϕ packaged into a single partial term using disjunction of cofibrations:

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma \vdash r, s : \mathbb{I} \quad \Gamma \vdash \phi \text{ cof} \quad \Gamma.\mathbb{I}.(a = r[\mathbb{p}] \vee \phi[\mathbb{p}]) \vdash a : A[\mathbb{p}^2]}{\Gamma \vdash \mathbf{hcomp}_A^{r \rightarrow s}(\phi, a) : A}$$

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma \vdash r, s : \mathbb{I} \quad \Gamma \vdash \phi \text{ cof} \quad \Gamma \vdash \phi \text{ true} \quad \Gamma.\mathbb{I}.(a = r[\mathbb{p}] \vee \phi[\mathbb{p}]) \vdash a : A[\mathbb{p}^2]}{\Gamma \vdash \mathbf{hcomp}_A^{r \rightarrow s}(\phi, a) = a[\mathbf{id}.s.\star] : A}$$

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma \vdash r : \mathbb{I} \quad \Gamma \vdash \phi \text{ cof} \quad \Gamma.\mathbb{I}.(a = r[\mathbb{p}] \vee \phi[\mathbb{p}]) \vdash a : A[\mathbb{p}^2]}{\Gamma \vdash \mathbf{hcomp}_A^{r \rightarrow r}(\phi, a) = a[\mathbf{id}.r.\star] : A}$$

With **hcomp** to hand, we were able to complete the necessary “programming exercise” implementing **coe** in **Path**. In so doing, however, we have unleashed another avalanche of necessary programming exercises: we must discuss how **hcomp** can be reduced for each type constructor. Fortunately, however, at this point we have all the necessary tools to do this for every connective except the universe. We discuss the rules governing **hcomp** for the non-universe connectives in Section 5.5, but they are largely unsurprising.

The real complexity of **hcomp** comes in defining $\mathbf{hcomp}_U^{r \rightarrow s}(\phi, A)$. The problem is that, as an element of the universe, this composition is a code for a type and so one must describe the type $\mathbf{El}(\mathbf{hcomp}_U^{r \rightarrow s}(\phi, A))$. It not obvious, but the constraints of **hcomp** mean that this type must be non-empty and so non-trivial introduction and elimination rules must be given to govern this type. As with any other type we must describe also the behavior of **hcomp** and **coe** in $\mathbf{El}(\mathbf{hcomp}_U^{r \rightarrow s}(\phi, A))$ and these “nested” composition problems are rather intricate.

This complexity, however, is the essential tool by which cubical type theory supports a computational account of univalence. We will return to this topic in Section 5.5, so we provide only the intuition here. Recall that the univalence axiom provides an inverse to a certain map $\mathbf{Path}(U, A, B) \rightarrow \mathbf{El}(A) \simeq \mathbf{El}(B)$. The domain of this map now consists of certain lines in the universe—codes of types depending on \mathbb{I} —and so to interpret univalence,

it suffices to define a family of types depending $A, B : \mathbf{U}$, an equivalence $e : \mathbf{El}(A) \simeq \mathbf{El}(B)$ and a dimension term $r : \mathbb{I}$. This type is typically written $\mathbf{V}(r, A, B, e)$ as in, univalence.

The idea is that this type must collapse to A when the interval variable is specialized to 0 and to B when it is specialized with 1. This type, by definition, is a path in the universe $\mathbf{Path}(\mathbf{U}, A, B)$. As with any other type, one must describe composition and coercion in this line of types and it is here that the invertability of the given map $\mathbf{El}(A) \simeq \mathbf{El}(B)$ is crucial: it is this map which is used to supply coercions from one end of \mathbf{V} to the other.

While this sketch omits a great many details—even simplifying the shape of \mathbf{V} slightly—this is the crucial idea and payoff for recasting identity types as path types. By forcing identity types in the universe to take this more flexible form, we can define novel type formers which themselves implement the novel identifications mandated by univalence. The details vary greatly between presentations, but this general strategy is ubiquitous: (1) using an interval to encode identity types as path types, (2) adding additional operations to all types to force these path types to be symmetric, transitive, *etc.* and (3) implementing univalence by a specific type family depending on the interval.

The reward for the complexity of cubical type theory is the following theorem.

Theorem 5.4.13. *Cubical type theory enjoys consistency, canonicity, and normalization.*

These theorems were established over several years, for several different variations of cubical type theory. The consistency of the theory was proven in the first papers on cubical type theory [CCHM18; AFH17]. Canonicity was established by Huber [Hub18] and Angiuli, Hou (Favonia), and Harper [AFH17]. Normalization was proven by Sterling and Angiuli [SA21].

5.5★ *Computing with coercions and compositions*

Section 5.4 presented the core aspects of cubical type theory, but with many rules and details elided. In this section, we endeavor to fill in a few of these gaps by explaining some of the rules governing the computation of **hcomp** and **coe** in various types. Even in a dedicated section, however, we will not provide all of these rules. A complete set can be found in *e.g.*, Angiuli et al. [Ang+21].

Fortunately, the remaining rules do not introduce new judgmental structure. Instead, they are more akin to programming exercises and show how to build *e.g.*, $\mathbf{hcomp}_{\Pi(A,B)}$ in terms of composition and coercion in A and B . Accordingly, while the previous section was replete with rules and substitutions, we shall see far fewer of these in this section. Instead, we shall focus on these “programming exercises” and often write out the resulting terms for computing composition and coercion in more informal type-theoretic notation. We will present a few examples for how these are turned into actual formal rules to be added to cubical type theory but thereafter leave this mechanical task to the reader.

Notation 5.5.1. In order to facilitate writing informal terms with `coe` and `hcomp`, we shall treat them as closed elements of the following types:

$$\begin{aligned} \mathbf{coe} &: (A : \mathbb{I} \rightarrow \mathbf{U})(i, j : \mathbb{I}) \rightarrow A(i) \rightarrow A(j) \\ \mathbf{hcomp}_\phi &: (A : \mathbf{U})(i, j : \mathbb{I})(a : (k : \mathbb{I}) \rightarrow (i = k \vee \phi) \rightarrow A) \rightarrow A \end{aligned}$$

5.5.1 `coe` for Π and Σ

We begin by describing coercion for dependent products and sums. These two examples contain all the interesting structure one finds in the definitions of `coe` for the types of base Martin-Löf type theory and so we give them a fair bit of attention.

We begin by specifying the right-hand side of the following definitional equality:

$$\frac{\Gamma.\mathbb{I} \vdash A \text{ type} \quad \Gamma.\mathbb{I}.A \vdash B \text{ type} \quad \Gamma \vdash r, s : \mathbb{I} \quad \Gamma \vdash p : \Sigma(A, B)[\mathbf{id}.r]}{\Gamma \vdash \mathbf{coe}_{\Sigma(A, B)}^{r \rightarrow s}(p) = ? : \Sigma(A, B)[\mathbf{id}.s]} \quad \text{e}$$

This is one of the many, many “programming exercises” in cubical type theory. Our goal shall be to produce a term using `coe` for A and B which has the appropriate type to fit into the above rule, subject to the additional condition that when $r = s$ then this term is equal to p . This last point is not strictly necessary for the rule to be well-formed, but it is an important sanity check. After all, the definitional equality for $\mathbf{coe}_{\Sigma(A, B)}^{r \rightarrow s}(p)$ will force this to be true by transitivity and so it makes sense to ensure that this forced equality is sensible.

We shall divide this process up into two steps. First, we present this term using informal type theory and second, we shall list out the formal term in proper notation.

Lemma 5.5.2. Fix $A : \mathbb{I} \rightarrow \mathbf{U}$, $B : (i : \mathbb{I}) \rightarrow A(i) \rightarrow \mathbf{U}$, $r, s : \mathbb{I}$, and $p : \sum_{a:A r} B r a$. Using `coe` for A and B , we can construct type $\mathbf{coe}(\lambda i \rightarrow \sum_{a:A i} B i a) r s p : \sum_{a:A s} B s a$ which is definitional equal to p if $r = s$.

Proof. By the η law for dependent sums, this term must be of the form (a, b) for some element of $A s$ and of $B s a$. In fact, it is straightforward to find the first component of this pair: $a = \mathbf{coe} A r s \mathbf{fst}(p)$.

The second component of the pair is more complex. Naïvely, one might hope that one could mirror the construction for a and use `coe` B in some manner. However, this is not well-typed! After all, B is not in the correct shape for `coe`: it is an element of $(i : \mathbb{I}) \rightarrow A i \rightarrow \mathbf{U}$ and not the required $\mathbb{I} \rightarrow \mathbf{U}$. Accordingly, to apply `coe` we must choose some element of A with which to specialize B . In fact, the situation is more fraught than this: A itself depends on \mathbb{I} and so if we wish to obtain a specialization of B with the type $\mathbb{I} \rightarrow \mathbf{U}$, we will require an element of $\bar{a} : (i : \mathbb{I}) \rightarrow A i$. Given such an \bar{a} , however, We can then use `coe` with $B_{\bar{a}} = \lambda i \rightarrow B i (a i)$ to attempt to construct b .

We can further narrow things down with this in mind. After all, our goal is to set $b = \mathbf{coe} B_{\bar{a}} r s \mathbf{snd}(p)$ and if this is to be type-correct we must have $\bar{a} r = \mathbf{fst}(p)$. Moreover, since we wish to have $b : B a s$ and so $\bar{a} s$ must be $a = \mathbf{coe} A r s \mathbf{fst}(p)$.

In order to obtain \bar{a} , we take advantage of the flexibility of \mathbf{coe} to coerce from r to a variable dimension, rather than 0 or 1. Specifically, we define \bar{a} as follows:

$$\bar{a} := \lambda i \rightarrow \mathbf{coe} A r i a$$

With \bar{a} to hand, we choose $b := \mathbf{coe} B_{\bar{a}} r s \mathbf{snd}(p)$, completing the required term. We leave it to the reader to check the required definitional equality holds when $r = s$. \square

Rendering the above term in formal notation, the rule can be completed to the following:

$$\frac{\Gamma.\mathbb{I} \vdash A \text{ type} \quad \Gamma.\mathbb{I}.A \vdash B \text{ type} \quad \Gamma \vdash r, s : \mathbb{I} \quad \Gamma \vdash p : \Sigma(A, B)[\mathbf{id}.r]}{\Gamma \vdash \mathbf{coe}_{\Sigma(A, B)}^{r \rightarrow s}(p) = \mathbf{pair}(\mathbf{coe}_A^{r \rightarrow s}(\mathbf{fst}(p)), \mathbf{coe}_{B[\mathbf{id}. \mathbf{coe}_{A[\mathbf{p}.i]}^{r \rightarrow q}(a[\mathbf{p}]]]}^{r \rightarrow s}(\mathbf{snd}(p))) : \Sigma(A, B)[\mathbf{id}.s]}$$

While the translation is largely mechanical, the reader can hopefully appreciate that the informal term is far more legible than the formal cousin!

We now turn to the case of dependent products. The process is mostly similar and we use the coercion operations on A and B to specify how $\mathbf{coe}_{\Pi(A, B)}$ ought to compute. Our goal is once more to fill in the right-hand side of the following equality:

$$\frac{\Gamma.\mathbb{I} \vdash A \text{ type} \quad \Gamma.\mathbb{I}.A \vdash B \text{ type} \quad \Gamma \vdash r, s : \mathbb{I} \quad \Gamma \vdash f : \Pi(A, B)[\mathbf{id}.r]}{\Gamma \vdash \mathbf{coe}_{\Pi(A, B)}^{r \rightarrow s}(f) = ? : \Pi(A, B)[\mathbf{id}.s]}$$

Lemma 5.5.3. Fix $A : \mathbb{I} \rightarrow \mathbf{U}$, $B : (i : \mathbb{I}) \rightarrow A(i) \rightarrow \mathbf{U}$, $r, s : \mathbb{I}$, and $p : (a : A r) \rightarrow B r a$. Using \mathbf{coe} for A and B , we can construct type $\mathbf{coe}(\lambda i \rightarrow (a : A i) \rightarrow B i a) r s p : (a : A s) \rightarrow B s a$ which is definitional equal to p if $r = s$.

Proof. Our goal is to construct an element of $(a : A(s)) \rightarrow B s a$ and, accordingly, we fix $a : A(s)$ and set about constructing $B s a$. We begin by defining $a_r = \mathbf{coe} A s r a$ such that we obtain $b_r = f(a_r) : B r a_r$. We would like to coerce b_r to obtain our desired element of $B s a$, but along what type should this coercion occur? We must find some $\bar{a} : (i : \mathbb{I}) \rightarrow A(i)$ such that $\bar{a}(r) = \mathbf{coe} A s r a$ and $\bar{a}(s) = a$. Capitalizing on the fact that $\mathbf{coe} A s s a = a$, we choose \bar{a} to be $\lambda k \rightarrow \mathbf{coe} A s k a$. The full term then becomes the following:

$$\lambda a \rightarrow \mathbf{coe}(\lambda k \rightarrow B k(\mathbf{coe} A s k a)) r s(f(\mathbf{coe} A s r a))$$

Once again, we leave it to the intrepid reader to confirm that if $r = s$ then this term is simply equivalent to f . \square

For the final time, we provide a translation of this informal definition into formal notation. Hereafter, we shall leave this mechanical (if tedious) process to the reader:

$$\frac{\Gamma.\mathbb{I} \vdash A \text{ type} \quad \Gamma.\mathbb{I}.A \vdash B \text{ type} \quad \Gamma \vdash r, s : \mathbb{I} \quad \Gamma \vdash f : \Pi(A, B)[\mathbf{id}.r]}{\Gamma \vdash \mathbf{coe}_{\Pi(A, B)}^{r \rightarrow s}(f) = \lambda(\mathbf{coe}_{B[p \circ q, \mathbf{coe}_{A[(p \circ p) \circ \mathbb{I}]}^{s[p] \rightarrow \mathbb{I}}]}(q[p])}^{r[p] \rightarrow s[p]}(\mathbf{app}(f[p], \mathbf{coe}_{A[p \circ q]}^{s[p] \rightarrow r[p]}(q)))) : \Pi(A, B)[\mathbf{id}.s]}$$

Undeniably, these rules are complicated.⁸ They are, however, really just a sequence of programming exercises and share many characteristics and so describing the first few cases is the most painful.

The next novelty, as already mentioned, comes in the definition of **coe** for path types. We turn to this next and, consequently, shift our attention to the second operator we must define for every type: **hcomp**.

5.5.2 Working with homogeneous composition operator

A common challenge when one begins to study cubical type theory is to “visualize” **hcomp**. While **coe** matched closely enough with the already familiar **subst** operator, the homogeneous composition operator is quite different than any of the combinators one typically encounters in intensional type theory. Prior to using it to compute coercion in path types, we give a few simple worked examples of **hcomp** to help demystify this operator.

Composing two paths using hcomp Let us begin cultivating intuition for **hcomp** by showing how we can use it to compose two paths $p_1 : \mathbf{Path}(A, a, b)$ and $p_2 : \mathbf{Path}(A, b, c)$. We shall do this using **hcomp** A , so it remains only to choose (1) the $r \rightarrow s$ direction we wish to compose along and (2) the cofibration ϕ to restrict along.

To visualize this situation, let us briefly fix two dimension variables $i, j : \mathbb{I}$ and instantiate p_1 with i and p_2 with j . We can draw the situation as follows:

$$\begin{array}{ccc} & & i \rightarrow \\ j \downarrow & & \\ & a & \xrightarrow{p_1(i)} & b \\ & \downarrow a & & \downarrow p_2(j) \\ & a & & c \end{array}$$

In anticipation of what is to come, we have added a “degenerate” edge corresponding to reflexivity along a . In order to construct the composite of our two edges, it suffices to

⁸Indeed, the reader may wonder how the authors managed to get these complicated terms correct. The answer is simple: they did not. They found numerous typos in the process of editing this section.

find a line which joins the bottom two vertices. It is here we invoke **hcomp**. Since our goal is to fill “down” in the j direction, *e.g.*, to push the top edge along the two vertical edges, we shall apply **hcomp** from 0 to 1. The cofibration shall be used to isolate the two sides in this direction we possess so $\phi := i = 0 \vee i = 1$.

Let us put these pieces of intuition together into a term. Our goal is to construct a path in A , so we will begin by binding a dimension variable $i : \mathbb{I}$. We then define the composite path as follows:

$$(p_2 \bullet p_1) i = \mathbf{hcomp}_\phi A 0 1 (\lambda k, _ \rightarrow [k = 0 \hookrightarrow p_1 i \mid \phi \hookrightarrow [i = 0 \hookrightarrow a \mid i = 1 \hookrightarrow p_2 k]])$$

Exercise 5.17. Argue that $p_2 \bullet p_1$ has the expected boundary *i.e.* that $(p_2 \bullet p_1) 0 = a$ and that $(p_2 \bullet p_1) 1 = c$.

What if we wish to obtain not just the bottom edge of the square, but the entire 2-dimensional term? Just as we could produce lines by using **coe** with a variable dimension as the target, we can “**hcomp** to the middle” using a dimension variable to obtain the entire square. We represent this with the following diagram:

$$\begin{array}{ccc}
 & \begin{array}{c} i \\ \rightarrow \\ j \downarrow \end{array} & \\
 a & \xrightarrow{p_1(i)} & b \\
 \downarrow a & \mathbf{hcomp}_\phi A 0 j \lambda k, _ \rightarrow \left[\begin{array}{l} k = 0 \hookrightarrow p_1 i \\ i = 0 \hookrightarrow a \\ i = 1 \hookrightarrow p_2 k \end{array} \right] & \downarrow p_2(j) \\
 a & \xrightarrow{(p_2 \bullet p_1) i} & c
 \end{array}$$

Exercise 5.18. Check that this 2-dimensional term has the relevant boundary conditions. In particular, if $j = 0$ check it collapses to $p_1 i$.

Inverting a path using hcomp For a second example, suppose we are given $p : \mathbf{Path}(A, a, b)$. We show how **hcomp** may be used to construct an inverse path $\mathbf{Path}(A, b, a)$. Once more, we shall fill a square involving p alongside two degenerate paths.

To visualize this situation, let us fix $i, j : \mathbb{I}$ and consider the following three lines:

$$\begin{array}{ccc}
 & & i \\
 & & \rightarrow \\
 j & \downarrow & \\
 & & a \xrightarrow{\quad a \quad} a \\
 & & \downarrow \quad \downarrow \\
 & & p(i) \quad a \\
 & & \downarrow \quad \downarrow \\
 & & b \quad a
 \end{array}$$

In order to compose paths, we have already shown how to use **hcomp** to complete these three edges to a square. The same general procedure applies, though the result is now the inverse to p . In particular, we have the following:

$$p^{-1} i := \mathbf{hcomp}_\phi A 0 1 \lambda k, _ \rightarrow \left[\begin{array}{l} k = 0 \hookrightarrow a \\ i = 0 \hookrightarrow p(i) \\ i = 1 \hookrightarrow a \end{array} \right]$$

In fact, with further effort we could use **hcomp** to construct higher paths witnessing that e.g., a path between $p \bullet p^{-1}$ and a constant path. Rather than pursuing this more fully, however, we return to the original example which prompted this detour.

Coercion in path types from hcomp We can now complete the loop that motivated this detour and show how to implement coercion in **Path**. Crucially, this requires both **coe** and **hcomp** working in concert.

Lemma 5.5.4. Fix $A : \mathbb{I} \rightarrow \mathbf{U}$, $a, b : (i : \mathbb{I}) \rightarrow A(i)$ alongside $r, s : \mathbb{I}$ and $p : \mathbf{Path}(A(r), a(r), b(r))$. Using **hcomp** and **coe** for A , there exists a term $\mathbf{coe} A r s p : \mathbf{Path}(A(s), a(s), b(s))$ which is definitionally equal to p when $r = s$.

Proof. As before, let us fix $k : \mathbb{I}$ such that it now suffices to define an element of $A(s)$ which specializes to $a(s)$ and $b(s)$ when $k = 0$ or $k = 1$. This latter condition is the sort of problem well-addressed by **hcomp** $_\phi$ where $\phi := k = 0 \vee k = 1$: one of the definitional equalities governing the construction precisely allows us to guarantee these equations. It remains to work out the direction in which we ought to apply **hcomp** $_\phi$ as well as the “top” of the square we are filling. Let us revisit the drawing of the situation we encountered

when first attempting to construct `coe` in `Path`:

$$\begin{array}{ccc}
 & \xrightarrow{k} & \\
 & & \text{coe } A r s (p k) \\
 \text{coe } A r s (a r) & \xrightarrow{\quad} & \text{coe } A r s (b r) \\
 \text{\scriptsize \textit{a s}} & & \text{\scriptsize \textit{b s}}
 \end{array}$$

Here, we have depicted the vertical lines as “wavy” since they do not actually form a path with the top corresponding to 0 and the bottom to 1. Instead, they represent lines in $A(s)$ such that e.g., when specialized with r become $\text{coe } A r s (b r)$ and at s become $b s$. This, however, is precisely what we require if we use composition from r to s , rather than from 0 to 1. All told then, we arrive at the following term:

$$\text{coe } A r s p := \lambda i \rightarrow \mathbf{hcomp}_{i=0 \vee i=1} (A s) \lambda k, _ \rightarrow \left[\begin{array}{l} k = r \hookrightarrow \text{coe } A r s (p i) \\ i = 0 \hookrightarrow \text{coe } A k s (a k) \\ i = 1 \hookrightarrow \text{coe } A k s (b k) \end{array} \right]$$

We leave it to the reader to confirm that all three of the branches of the disjunction match as required on their overlaps and that when $r = s$ this term collapses to p . \square

5.5.3 Unfolding `hcomp` in various type constructors

While we have discussed the core rules governing `coe` at this point, it remains to do so for `hcomp`. Just as with coercion, for specifying core connectives amounts to a sequence of programming exercises and we give the details only for dependent products and path types.

Lemma 5.5.5. *Fix a cofibration ϕ , types $A : \mathbf{U}$, $B : A \rightarrow \mathbf{U}$, dimension terms $r, s : \mathbb{I}$, and a partial term $f : (i : \mathbb{I}) \rightarrow (i = r \vee \phi) \rightarrow (a : A) \rightarrow B(a)$. There exists a term $\mathbf{hcomp}_\phi ((a : A) \rightarrow B a) r s a : A$ built from composition in B satisfying the expected definitional equalities.*

Proof. Let us fix $a : A$ such that we must build $b : B(a)$ such that (1) if $r = s$ then $b = f s a$ and if (2) ϕ holds then $b = f s a$. To this end, we shall use composition in $B(a)$:

$$b = \mathbf{hcomp}_\phi (B a) r s \lambda i, _ \rightarrow f i _ a$$

It is routine to see that this gives rise to the required term using the boundary conditions of $\mathbf{hcomp}_\phi (B a) r s$. \square

Lemma 5.5.6. Fix a cofibration ϕ , a type $A : \mathbf{U}$, elements $a, b : A$, dimension terms $r, s : \mathbb{I}$, and a partial term $p : (i : \mathbb{I}) \rightarrow (i = r \vee \phi) \rightarrow \mathbf{Path}(A, a, b)$. There exists a term $\mathbf{hcomp}_\phi(\mathbf{Path}(A, a, b)) r s p : \mathbf{Path}(A, a, b)$ built from composition and coercion in A satisfying the expected definitional equalities.

Proof. The required term is an application of \mathbf{hcomp} in A . Since we intend to construct a path, we fix $i : \mathbb{I}$ such that $\lambda j \rightarrow p j _ i : (j : \mathbb{I}) \rightarrow j = r \vee \phi \rightarrow A$ is a partial element suitable as input for \mathbf{hcomp} .

This is almost sufficient, but we must also ensure that the resulting extended term satisfies the boundary condition necessary to form an element of $\mathbf{Path}(A, a, b)$. To fix these boundaries, we extend ϕ with faces to govern the behavior of this term when $i = 0$ or $i = 1$. The final term is given as follows:

$$\mathbf{hcomp}_\phi(\mathbf{Path}(A, a, b)) r s p := \lambda i \rightarrow \mathbf{hcomp}_{\phi \vee_{i=0} \vee_{i=1}} A r s \lambda j, _ \rightarrow \left[\begin{array}{l} \phi \vee j = r \hookrightarrow p j _ i \\ i = 0 \hookrightarrow a \\ i = 1 \hookrightarrow b \end{array} \right]$$

We once more leave it to the reader to check that this satisfies the necessary boundary conditions. \square

5.5.4 \mathbf{V} and univalence

Finally, we turn to the rules necessary to animate both \mathbf{hcomp} in \mathbf{U} and univalence. The crucial idea behind both the same: paths in the universe are, by definition, codes which depend on \mathbb{I} and so to implement either \mathbf{hcomp} or univalence, it suffices to define new types. We shall focus largely on the new type necessary to implement univalence \mathbf{V} , but much the process transfers to \mathbf{hcomp} .

Suppose we are given $A, B : \mathbf{U}$ along with $e : A \simeq B$. We wish to construct a path $ua e : \mathbf{Path}(\mathbf{U}, A, B)$ or, equivalently, a map $p : \mathbb{I} \rightarrow \mathbf{U}$ such that $p 0 = A$ and $p 1 = B$. We intend for ua to be inverse to the canonical map $\mathbf{idToEquiv} : \mathbf{Id}(\mathbf{U}, A, B) \rightarrow A \simeq B$ which, in this setting, amounts to the requiring that $\mathbf{coe} p 0 1 : A \rightarrow B$ is equal to e .

Remark 5.5.7. The reader may wonder whether we need an additional constraint ensuring that $ua(\mathbf{coe} p 0 1)$ can be identified with p . As we remarked in Section 5.2, this direction holds automatically. \diamond

Our goal shall be to define a new type $\mathbf{V}(A, B, e, r)$ and to set $ua e := \lambda i \rightarrow \mathbf{V}(A, B, e, i)$. We begin with the (provisional) formation rule for \mathbf{V} :

$$\frac{\Gamma \vdash A, B \text{ type} \quad \Gamma \vdash e : A \simeq B \quad \Gamma \vdash r : \mathbb{I}}{\Gamma \vdash \mathbf{V}(A, B, e, r) \text{ type}}$$

We note that this definition is sensitive to the precise realization of equivalence we choose. However, any of the notions presented in Section 5.2 suffice and so we shall ignore this detail. Moreover, we must ensure that our universe is closed under \mathbf{V} in order to actually carry out the definition of ua . It is more convenient to specify rules for the type \mathbf{V} rather than the code, however, and so we shall focus on that.

The above set of constraints on ua and path types generally translate into the following requirements for $\mathbf{V}(A, B, e, r)$:

- We must have definitional equalities $\mathbf{V}(A, B, e, 0) = A$ and $\mathbf{V}(A, B, e, 1) = B$.
- It must be the case that $\mathbf{coe}(\lambda i \rightarrow \mathbf{V}(A, B, e, i)) 0 1 = e$.
- We must be able to implement \mathbf{hcomp} and \mathbf{CoeSym} for \mathbf{V} .

Of course, $\lambda i \rightarrow \mathbf{V}(A, B, e, i)$ is always fully constrained up to equivalence: it is the unique inhabitant of $\mathbf{Path}(\mathbf{U}, A, B)$ sent to e by \mathbf{coe} . In this way, it is largely unimportant how precisely \mathbf{V} is realized. What matters is only that such a type can exist and satisfy the list of required properties. To this end, these constraints are useful for nailing down the particular rules which define \mathbf{V} more precisely and, unfortunately, we must give new rules. \mathbf{V} cannot be defined by a clever combination of existing type formers because of the first requirement; we presently have no means of defining a type which degenerates to two distinct types depending on the endpoints of an interval.

In fact, given all these constraints there are precious few valid choices for the introduction and elimination rules of \mathbf{V} . The difficulty is that it is not obvious whether any given choice of rules will suffice until one carefully checks each condition. Accordingly, we will present the correct rules below and only then discuss some of the subtleties:

$$\frac{\Gamma \vdash A, B \text{ type} \quad \Gamma \vdash e : A \simeq B}{\Gamma \vdash \mathbf{V}(A, B, e, 0) = A \text{ type} \quad \Gamma \vdash \mathbf{V}(A, B, e, 1) = B \text{ type}}$$

$$\frac{\Gamma \vdash A, B \text{ type} \quad \Gamma \vdash e : A \simeq B \quad \Gamma \vdash r : \mathbb{I} \quad \Gamma, r = 0 \vdash a : A[\mathbb{p}] \quad \Gamma \vdash b : B \quad \Gamma, r = 0 \vdash \mathbf{app}(e[\mathbb{p}], a) = b[\mathbb{p}] : B[\mathbb{p}]}{\Gamma \vdash \mathbf{Vin}(a, b, r) : \mathbf{V}(A, B, e, r)}$$

$$\Gamma, r = 0 \vdash \mathbf{Vin}(a, b, r)[\mathbb{p}] = a : A[\mathbb{p}] \quad \Gamma, r = 1 \vdash \mathbf{Vin}(a, b, r)[\mathbb{p}] = b : B[\mathbb{p}]$$

$$\frac{\Gamma \vdash A, B \text{ type} \quad \Gamma \vdash e : A \simeq B \quad \Gamma \vdash r : \mathbb{I} \quad \Gamma \vdash v : \mathbf{V}(A, B, e, r)}{\Gamma \vdash \mathbf{Vout}(v) : B}$$

$$\Gamma, r = 0 \vdash \mathbf{Vout}(v) = \mathbf{app}(e, v) : B[\mathbb{p}] \quad \Gamma, r = 1 \vdash \mathbf{Vout}(v) = v : B[\mathbb{p}]$$

$$\begin{array}{c}
\Gamma \vdash A, B \text{ type} \quad \Gamma \vdash e : A \simeq B \quad \Gamma \vdash r : \mathbb{I} \\
\Gamma, r = 0 \vdash a : A[\rho] \quad \Gamma \vdash b : B \quad \Gamma, r = 0 \vdash \mathbf{app}(e[\rho], a) = b[\rho] : B[\rho] \\
\hline
\Gamma \vdash \mathbf{Vout}(\mathbf{Vin}(a, b, r)) = b : B \\
\\
\Gamma \vdash A, B \text{ type} \quad \Gamma \vdash e : A \simeq B \quad \Gamma \vdash r : \mathbb{I} \quad \Gamma \vdash v : \mathbf{V}(A, B, e, r) \\
\hline
\Gamma \vdash \mathbf{Vin}(v[\rho], \mathbf{Vout}(v), r) = v : \mathbf{V}(A, B, e, r)
\end{array}$$

In total then, an element of $\mathbf{V}(A, B, e, r)$ contains a partial element of A and a full element of B which *match up* according to e when both are defined. The introduction and elimination rules (along with their β and η principles) are then nearly routine from this perspective. The complexity comes from the various rules which apply if $r = 0$ or $r = 1$.

These are a consequence of having $\mathbf{V}(A, B, e, r)$ collapse definitionally to A and B . We have not encountered rules similar to this with other type formers and they impose a number of unique constraints on the rules around \mathbf{V} if we are to avoid having terms of $\mathbf{V}(A, B, e, r)$ polluting A and B . For instance, we must add rules ensuring that $\mathbf{Vin}(a, b)$ correctly equates to a or b where this is required. Similarly, $\mathbf{Vout}(v)$ cannot come only with a β rule to govern its behavior, as it must account for the situations where v becomes an ordinary element of A and B .

To illustrate the delicacy of these rules, imagine a simple possible replacement: instead of requiring $\Gamma, r = 0 \vdash \mathbf{app}(e[\rho], a) = b[\rho] : B[\rho]$, what if we required that $\mathbf{app}(e^{-1}, b)$ was definitionally equal to a ? While this is seemingly innocuous, e and e^{-1} are inverses only up to a path and not necessarily definitionally inverses. Consequently, this exchange would make it impossible to properly specify the behavior of $\mathbf{Vout}(v)$ when $r = 0$; depending on the order in which rules were applied one could obtain distinct (but path equal!) terms.

Another mysterious aspect of these rules is the asymmetry between A and B . Why a is required to be a partial element whereas b is total as opposed to defined only when $r = 1$ holds. What matters is not so much whether a or b is partial, but merely that one of the two is fully defined and one is not. If neither is fully defined, it becomes impossible to state that a and b are equated by e . More subtly, if both are fully defined it becomes impossible to specify \mathbf{coe} in \mathbf{V} .

The definitions of \mathbf{hcomp} and \mathbf{coe} in \mathbf{V} are complex and we will not attempt to detail them here. The interested reader should consult Appendix B of Angiuli [Ang19] for precise account of \mathbf{V} .

Finally, we note that the same chain of reasoning that leads to this definition of \mathbf{V} can be used to produce the type implementing $\mathbf{hcomp}_U^{r \rightarrow s}(\phi, A)$. We can once more list out the various definitional equalities which such a type must satisfy as well as what types it must be equivalent to. Unfurling these, we determine that elements of $\mathbf{hcomp}_U^{r \rightarrow s}(\phi, A)$ are essentially smaller formal composition problems, just as elements of \mathbf{V} were “suspended coercions along e ”. Unfortunately, the details and bookkeeping around

such formal composition problems (and composition problems *of* formal composition problems) is taxing. A curious reader should once again consult Angiuli [[Ang19](#)].

Martin-Löf type theory

This appendix presents a substitution calculus [ML92; Tas93; Dyb96] for several variants of Martin-Löf's dependent type theory. Martin-Löf type theories are systems admitting the rules in section *Contexts and substitutions*; the rules specific to extensional type theory, those axiomatizing extensional equality types, are marked (ETT); the rules specific to intensional type theory, those axiomatizing *intensional equality types*, are marked (ITT).

Judgments

Martin-Löf type theory has four basic judgments:

1. $\vdash \Gamma \text{ cx}$ asserts that Γ is a context.
2. $\Delta \vdash \gamma : \Gamma$, presupposing $\vdash \Delta \text{ cx}$ and $\vdash \Gamma \text{ cx}$, asserts that γ is a substitution from Δ to Γ (*i.e.*, assigns a term in Δ to each variable in Γ).
3. $\Gamma \vdash A \text{ type}$, presupposing $\vdash \Gamma \text{ cx}$, asserts that A is a type in context Γ .
4. $\Gamma \vdash a : A$, presupposing $\vdash \Gamma \text{ cx}$ and $\Gamma \vdash A \text{ type}$, asserts that a is an element/term of type A in context Γ .

The *presuppositions* of a judgment are its meta-implicit-arguments, so to speak. For instance, the judgment $\Gamma \vdash A \text{ type}$ is sensible to write (is meta-well-typed) only when the judgment $\vdash \Gamma \text{ cx}$ holds. We adopt the convention that asserting the truth of a judgment implicitly asserts its well-formedness; thus asserting $\Gamma \vdash A \text{ type}$ also asserts $\vdash \Gamma \text{ cx}$.

As we assert the existence of various contexts, substitutions, types, and terms, we will simultaneously need to assert that some of these (already introduced) objects are equal to other (already introduced) objects of the same kind.

1. $\Delta \vdash \gamma = \gamma' : \Gamma$, presupposing $\Delta \vdash \gamma : \Gamma$ and $\Delta \vdash \gamma' : \Gamma$, asserts that γ, γ' are equal substitutions from Δ to Γ .
2. $\Gamma \vdash A = A' \text{ type}$, presupposing $\Gamma \vdash A \text{ type}$ and $\Gamma \vdash A' \text{ type}$, asserts that A, A' are equal types in context Γ .
3. $\Gamma \vdash a = a' : A$, presupposing $\Gamma \vdash a : A$ and $\Gamma \vdash a' : A$, asserts that a, a' are equal elements of type A in context Γ .

Two types (*resp.*, contexts, substitutions, terms) being equal has the force that it does in standard mathematics: any expression can be replaced silently by an equal expression without affecting the meaning or truth of the statement in which it appears. One important example of this principle is the “conversion rule” which states that if $\Gamma \vdash A = A'$ type and $\Gamma \vdash a : A$, then $\Gamma \vdash a : A'$.

In the rules that follow, some arguments of substitution, type, and term formers are typeset as gray subtitles; these are arguments that we will often omit because they can be inferred from context and are tedious and distracting to write.

Contexts and substitutions

$$\begin{array}{c}
\frac{}{\vdash \mathbf{1} \text{ cx}} \text{ CX/EMP} \qquad \frac{\vdash \Gamma \text{ cx} \quad \Gamma \vdash A \text{ type}}{\vdash \Gamma.A \text{ cx}} \text{ CX/EXT} \\
\hline
\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{id}_\Gamma : \Gamma} \text{ SB/ID} \qquad \frac{\Gamma_2 \vdash \gamma_1 : \Gamma_1 \quad \Gamma_1 \vdash \gamma_0 : \Gamma_0}{\Gamma_2 \vdash \gamma_0 \circ_{\Gamma_2, \Gamma_1, \Gamma_0} \gamma_1 : \Gamma_0} \text{ SB/COMP} \\
\hline
\frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{id}_\Gamma \circ \gamma = \gamma : \Gamma} \qquad \frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \gamma \circ \mathbf{id}_\Delta = \gamma : \Gamma} \qquad \frac{\Gamma_3 \vdash \gamma_2 : \Gamma_2 \quad \Gamma_2 \vdash \gamma_1 : \Gamma_1 \quad \Gamma_1 \vdash \gamma_0 : \Gamma_0}{\Gamma_3 \vdash \gamma_0 \circ (\gamma_1 \circ \gamma_2) = (\gamma_0 \circ \gamma_1) \circ \gamma_2 : \Gamma_0} \\
\hline
\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type}}{\Delta \vdash A[\gamma]_{\Delta, \Gamma} \text{ type}} \text{ TY/SB} \qquad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : A}{\Delta \vdash a[\gamma]_{\Delta, \Gamma} : A[\gamma]} \text{ TM/SB} \\
\frac{\Gamma \vdash A \text{ type}}{\Gamma \vdash A[\mathbf{id}_\Gamma] = A \text{ type}} \qquad \frac{\Gamma \vdash a : A}{\Gamma \vdash a[\mathbf{id}_\Gamma] = a : A} \\
\frac{\Gamma_2 \vdash \gamma_1 : \Gamma_1 \quad \Gamma_1 \vdash \gamma_0 : \Gamma_0 \quad \Gamma_0 \vdash A \text{ type}}{\Gamma_2 \vdash A[\gamma_0 \circ \gamma_1] = A[\gamma_0][\gamma_1] \text{ type}} \qquad \frac{\Gamma_2 \vdash \gamma_1 : \Gamma_1 \quad \Gamma_1 \vdash \gamma_0 : \Gamma_0 \quad \Gamma_0 \vdash a : A}{\Gamma_2 \vdash a[\gamma_0 \circ \gamma_1] = a[\gamma_0][\gamma_1] : A[\gamma_0 \circ \gamma_1]} \\
\hline
\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash !_\Gamma : \mathbf{1}} \text{ SB/EMP} \qquad \frac{\Gamma \vdash \delta : \mathbf{1}}{\Gamma \vdash !_\Gamma = \delta : \mathbf{1}} \\
\hline
\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type} \quad \Delta \vdash a : A[\gamma]}{\Delta \vdash \gamma \cdot_{\Delta, \Gamma, A} a : \Gamma.A} \text{ SB/EXT} \qquad \frac{\Gamma \vdash A \text{ type}}{\Gamma.A \vdash \mathbf{p}_{\Gamma, A} : \Gamma} \text{ SB/WK}
\end{array}$$

$$\frac{\Gamma \vdash A \text{ type}}{\Gamma.A \vdash \mathbf{q}_{\Gamma,A} : A[\mathbf{p}_{\Gamma,A}]} \text{VAR} \quad \frac{\Delta \vdash \gamma : \Gamma \quad \Delta \vdash a : A[\gamma]}{\Delta \vdash \mathbf{p}_{\Gamma,A} \circ_{\Gamma,A} (\gamma.a) = \gamma : \Gamma} \quad \frac{\Delta \vdash \gamma : \Gamma \quad \Delta \vdash a : A[\gamma]}{\Delta \vdash \mathbf{q}_{\Gamma,A}[\gamma.a] = a : A[\gamma]}$$

$$\frac{\Delta \vdash \gamma : \Gamma.A}{\Delta \vdash \gamma = (\mathbf{p}_{\Gamma,A} \circ_{\Gamma,A} \gamma).(\mathbf{q}_{\Gamma,A}[\gamma]) : \Gamma.A}$$

Π -types

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type}}{\Gamma \vdash \Pi_{\Gamma}(A, B) \text{ type}} \text{PI/FORM} \quad \frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash b : B}{\Gamma \vdash \lambda_{\Gamma,A,B}(b) : \Pi(A, B)} \text{PI/INTRO}$$

$$\frac{\Gamma \vdash a : A \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash f : \Pi(A, B)}{\Gamma \vdash \mathbf{app}_{\Gamma,A,B}(f, a) : B[\mathbf{id}_{\Gamma}.a]} \text{PI/ELIM}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type}}{\Delta \vdash \Pi_{\Gamma}(A, B)[\gamma] = \Pi_{\Delta}(A[\gamma], B[(\gamma \circ \mathbf{p}_{\Delta,A[\gamma]}) \cdot \mathbf{q}_{\Delta,A[\gamma]}]) \text{ type}}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type} \quad \Gamma.A \vdash b : B}{\Delta \vdash \lambda(b)[\gamma] = \lambda(b[(\gamma \circ \mathbf{p}) \cdot \mathbf{q}]) : \Pi(A, B)[\gamma]}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : A \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash f : \Pi(A, B)}{\Delta \vdash \mathbf{app}(f, a)[\gamma] = \mathbf{app}(f[\gamma], a[\gamma]) : B[(\mathbf{id}_{\Gamma}.a) \circ \gamma]}$$

$$\frac{\Gamma \vdash a : A \quad \Gamma.A \vdash b : B}{\Gamma \vdash \mathbf{app}(\lambda(b), a) = b[\mathbf{id}.a] : B[\mathbf{id}.a]} \quad \frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash f : \Pi(A, B)}{\Gamma \vdash f = \lambda(\mathbf{app}(f[\mathbf{p}_{\Gamma,A}], \mathbf{q}_{\Gamma,A})) : \Pi(A, B)}$$

Σ -types

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type}}{\Gamma \vdash \Sigma_{\Gamma}(A, B) \text{ type}} \text{SIGMA/FORM}$$

$$\frac{\Gamma \vdash a : A \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash b : B[\mathbf{id}_{\Gamma}.a]}{\Gamma \vdash \mathbf{pair}_{\Gamma,A,B}(a, b) : \Sigma(A, B)} \text{SIGMA/INTRO}$$

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash p : \Sigma(A, B)}{\Gamma \vdash \mathbf{fst}_{\Gamma, A, B}(p) : A} \text{ SIGMA/ELIM/FST}$$

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash p : \Sigma(A, B)}{\Gamma \vdash \mathbf{snd}_{\Gamma, A, B}(p) : B[\mathbf{id}_{\Gamma}.\mathbf{fst}(p)]} \text{ SIGMA/ELIM/SND}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type}}{\Delta \vdash \Sigma_{\Gamma}(A, B)[\gamma] = \Sigma_{\Delta}(A[\gamma], B[(\gamma \circ \mathbf{p}).\mathbf{q}]) \text{ type}}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : A \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash b : B[\mathbf{id}.a]}{\Delta \vdash \mathbf{pair}(a, b)[\gamma] = \mathbf{pair}(a[\gamma], b[\gamma]) : \Sigma(A, B)[\gamma]}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash p : \Sigma(A, B)}{\Delta \vdash \mathbf{fst}(p)[\gamma] = \mathbf{fst}(p[\gamma]) : A[\gamma]}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash p : \Sigma(A, B)}{\Delta \vdash \mathbf{snd}(p)[\gamma] = \mathbf{snd}(p[\gamma]) : B[(\mathbf{id}.\mathbf{fst}(p)) \circ \gamma]}$$

$$\frac{\Gamma \vdash a : A \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash b : B[\mathbf{id}.a]}{\Gamma \vdash \mathbf{fst}(\mathbf{pair}(a, b)) = a : A}$$

$$\frac{\Gamma \vdash a : A \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash b : B[\mathbf{id}.a]}{\Gamma \vdash \mathbf{snd}(\mathbf{pair}(a, b)) = b : B[\mathbf{id}.a]}$$

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash p : \Sigma(A, B)}{\Gamma \vdash p = \mathbf{pair}(\mathbf{fst}(p), \mathbf{snd}(p)) : \Sigma(A, B)}$$

Extensional equality types

$$\frac{\Gamma \vdash a : A \quad \Gamma \vdash b : A}{\Gamma \vdash \mathbf{Eq}_{\Gamma}(A, a, b) \text{ type}} \text{ EQ/FORM (ETT)} \quad \frac{\Gamma \vdash a : A}{\Gamma \vdash \mathbf{refl}_{\Gamma, A, a} : \mathbf{Eq}(A, a, a)} \text{ EQ/INTRO (ETT)}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : A \quad \Gamma \vdash b : A}{\Delta \vdash \mathbf{Eq}_{\Gamma}(A, a, b)[\gamma] = \mathbf{Eq}_{\Delta}(A[\gamma], a[\gamma], b[\gamma]) \text{ type}} \text{ (ETT)}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : A}{\Delta \vdash \mathbf{refl}[\gamma] = \mathbf{refl} : \mathbf{Eq}(A, a, a)[\gamma]} \text{ (ETT)}$$

$$\frac{\Gamma \vdash a : A \quad \Gamma \vdash b : A \quad \Gamma \vdash p : \mathbf{Eq}(A, a, b)}{\Gamma \vdash a = b : A} \text{ (ETT)}$$

$$\frac{\Gamma \vdash a : A \quad \Gamma \vdash b : A \quad \Gamma \vdash p : \mathbf{Eq}(A, a, b)}{\Gamma \vdash p = \mathbf{refl} : \mathbf{Eq}(A, a, b)} \text{ (ETT)}$$

Unit type

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{Unit}_\Gamma \text{ type}} \text{ UNIT/FORM}$$

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{tt}_\Gamma : \mathbf{Unit}} \text{ UNIT/INTRO}$$

$$\frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{Unit}_\Gamma[\gamma] = \mathbf{Unit}_\Delta \text{ type}}$$

$$\frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{tt}_\Gamma[\gamma] = \mathbf{tt}_\Delta : \mathbf{Unit}}$$

$$\frac{\Gamma \vdash a : \mathbf{Unit}}{\Gamma \vdash a = \mathbf{tt} : \mathbf{Unit}}$$

Empty type

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{Void}_\Gamma \text{ type}} \text{ EMPTY/FORM}$$

$$\frac{\Gamma \vdash b : \mathbf{Void} \quad \Gamma.\mathbf{Void} \vdash A \text{ type}}{\Gamma \vdash \mathbf{absurd}_{\Gamma, A}(b) : A[\mathbf{id}.b]} \text{ EMPTY/ELIM}$$

$$\frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{Void}_\Gamma[\gamma] = \mathbf{Void}_\Delta \text{ type}}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash b : \mathbf{Void} \quad \Gamma.\mathbf{Void} \vdash A \text{ type}}{\Delta \vdash \mathbf{absurd}(b)[\gamma] = \mathbf{absurd}(b[\gamma]) : A[\gamma.b[\gamma]]}$$

Boolean type

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{Bool}_\Gamma \text{ type}} \text{ BOOL/FORM}$$

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{true}_\Gamma : \mathbf{Bool}} \text{ BOOL/INTRO/TRUE}$$

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{false}_\Gamma : \mathbf{Bool}} \text{ BOOL/INTRO/FALSE}$$

$$\frac{\Gamma.\mathbf{Bool} \vdash A \text{ type} \quad \Gamma \vdash b : \mathbf{Bool} \quad \Gamma \vdash a_t : A[\mathbf{id}.true] \quad \Gamma \vdash a_f : A[\mathbf{id}.false]}{\Gamma \vdash \mathbf{if}_{\Gamma, A}(a_t, a_f, b) : A[\mathbf{id}.b]} \text{ BOOL/ELIM}$$

$$\frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{Bool}_\Gamma[\gamma] = \mathbf{Bool}_\Delta \text{ type}}$$

$$\frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{true}_\Gamma[\gamma] = \mathbf{true}_\Delta : \mathbf{Bool}} \quad \frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{false}_\Gamma[\gamma] = \mathbf{false}_\Delta : \mathbf{Bool}}$$

$$\frac{\Gamma \vdash b : \mathbf{Bool} \quad \Gamma.\mathbf{Bool} \vdash A \text{ type} \quad \Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a_t : A[\mathbf{id.true}] \quad \Gamma \vdash a_f : A[\mathbf{id.false}]}{\Delta \vdash \mathbf{if}(a_t, a_f, b)[\gamma] = \mathbf{if}(a_t[\gamma], a_f[\gamma], b[\gamma]) : A[\gamma.b[\gamma]]}$$

$$\frac{\vdash \Gamma \text{ cx} \quad \Gamma.\mathbf{Bool} \vdash A \text{ type} \quad \Gamma \vdash a_t : A[\mathbf{id.true}] \quad \Gamma \vdash a_f : A[\mathbf{id.false}]}{\Gamma \vdash \mathbf{if}(a_t, a_f, \mathbf{true}) = a_t : A[\mathbf{id.true}]}$$

$$\frac{\vdash \Gamma \text{ cx} \quad \Gamma.\mathbf{Bool} \vdash A \text{ type} \quad \Gamma \vdash a_t : A[\mathbf{id.true}] \quad \Gamma \vdash a_f : A[\mathbf{id.false}]}{\Gamma \vdash \mathbf{if}(a_t, a_f, \mathbf{false}) = a_f : A[\mathbf{id.false}]}$$

Natural number type

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{Nat}_\Gamma \text{ type}} \text{ NAT/FORM}$$

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{zero}_\Gamma : \mathbf{Nat}} \text{ NAT/INTRO/ZERO} \quad \frac{\Gamma \vdash n : \mathbf{Nat}}{\Gamma \vdash \mathbf{suc}_\Gamma(n) : \mathbf{Nat}} \text{ NAT/INTRO/SUC}$$

$$\frac{\Gamma.\mathbf{Nat} \vdash A \text{ type} \quad \Gamma \vdash a_z : A[\mathbf{id.zero}] \quad \Gamma.\mathbf{Nat}.A \vdash a_s : A[\mathbf{p}^2.\mathbf{suc}(q[p])] \quad \Gamma \vdash n : \mathbf{Nat}}{\Gamma \vdash \mathbf{rec}_{\Gamma,A}(a_z, a_s, n) : A[\mathbf{id}.n]} \text{ NAT/ELIM}$$

$$\frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{Nat}_\Gamma[\gamma] = \mathbf{Nat}_\Delta \text{ type}}$$

$$\frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{zero}_\Gamma[\gamma] = \mathbf{zero}_\Delta : \mathbf{Nat}} \quad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash n : \mathbf{Nat}}{\Delta \vdash \mathbf{suc}_\Gamma(n)[\gamma] = \mathbf{suc}_\Delta(n[\gamma]) : \mathbf{Nat}}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma.\mathbf{Nat} \vdash A \text{ type} \quad \Gamma \vdash a_z : A[\mathbf{id.zero}] \quad \Gamma.\mathbf{Nat}.A \vdash a_s : A[\mathbf{p}^2.\mathbf{suc}(q[p])] \quad \Gamma \vdash n : \mathbf{Nat}}{\Delta \vdash \mathbf{rec}(a_z, a_s, n)[\gamma] = \mathbf{rec}(a_z[\gamma], a_s[(\gamma \circ \mathbf{p}^2).q[p].q], n[\gamma]) : A[\gamma.n[\gamma]]}$$

$$\frac{\Gamma.\text{Nat} \vdash A \text{ type} \quad \Gamma \vdash a_z : A[\text{id.zero}] \quad \Gamma.\text{Nat}.A \vdash a_s : A[\text{p}^2.\text{suc}(q[p])]}{\Gamma \vdash \text{rec}(a_z, a_s, \text{zero}) = a_z : A[\text{id.zero}]}$$

$$\frac{\Gamma.\text{Nat} \vdash A \text{ type} \quad \Gamma \vdash a_z : A[\text{id.zero}] \quad \Gamma.\text{Nat}.A \vdash a_s : A[\text{p}^2.\text{suc}(q[p])] \quad \Gamma \vdash n : \text{Nat}}{\Gamma \vdash \text{rec}(a_z, a_s, \text{suc}(n)) = a_s[\text{id.n.rec}(a_z, a_s, n)] : A[\text{id.suc}(n)]}$$

Intensional equality types

$$\frac{\Gamma \vdash a : A \quad \Gamma \vdash b : A}{\Gamma \vdash \text{Id}_\Gamma(A, a, b) \text{ type}} \text{ID/FORM (ITT)} \quad \frac{\Gamma \vdash a : A}{\Gamma \vdash \text{refl}_{\Gamma, A, a} : \text{Id}(A, a, a)} \text{ID/INTRO (ITT)}$$

$$\frac{\Gamma \vdash a : A \quad \Gamma \vdash b : A \quad \Gamma \vdash p : \text{Id}(A, a, b) \quad \Gamma.A.A[p].\text{Id}(A[p^2], q[p], q) \vdash C \text{ type} \quad \Gamma.A \vdash c : C[p.q.q.refl]}{\Gamma \vdash \text{J}_{\Gamma, A, a, b, C}(c, p) : C[\text{id.a.b.p}]} \text{ID/ELIM (ITT)}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : A \quad \Gamma \vdash b : A}{\Delta \vdash \text{Id}_\Gamma(A, a, b)[\gamma] = \text{Id}_\Delta(A[\gamma], a[\gamma], b[\gamma]) \text{ type}} \text{(ITT)}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : A}{\Delta \vdash \text{refl}[\gamma] = \text{refl} : \text{Id}(A[\gamma], a[\gamma], a[\gamma])} \text{(ITT)}$$

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : A \quad \Gamma \vdash b : A \quad \Gamma \vdash p : \text{Id}(A, a, b) \quad \Gamma.A.A[p].\text{Id}(A[p^2], q[p], q) \vdash C \text{ type} \quad \Gamma.A \vdash c : C[p.q.q.refl]}{\Delta \vdash \text{J}(c, p)[\gamma] = \text{J}(c[(\gamma \circ p).q], p[\gamma]) : C[\gamma.a[\gamma].b[\gamma].p[\gamma]]} \text{(ITT)}$$

$$\frac{\Gamma \vdash a : A \quad \Gamma.A.A[p].\text{Id}(A[p^2], q[p], q) \vdash C \text{ type} \quad \Gamma.A \vdash c : C[p.q.q.refl]}{\Gamma \vdash \text{J}(c, \text{refl}) = c[\text{id.a}] : C[\text{id.a.a.refl}]} \text{(ITT)}$$

Universes

$$\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \text{U}_{\Gamma, i} \text{ type}} \text{UNI/FORM} \quad \frac{\Gamma \vdash a : \text{U}_i}{\Gamma \vdash \text{El}_{\Gamma, i}(a) \text{ type}} \text{EL/FORM}$$

$$\frac{\Gamma \vdash c_0 : \text{U}_i \quad \Gamma.\text{El}_i(c_0) \vdash c_1 : \text{U}_i}{\Gamma \vdash \text{pi}_{i, \Gamma}(c_0, c_1) : \text{U}_i} \text{PI/CODE} \quad \frac{\Gamma \vdash c_0 : \text{U}_i \quad \Gamma.\text{El}_i(c_0) \vdash c_1 : \text{U}_i}{\Gamma \vdash \text{sig}_{i, \Gamma}(c_0, c_1) : \text{U}_i} \text{SIG/CODE}$$

$$\begin{array}{c}
\frac{\Gamma \vdash c : \mathbf{U}_i \quad \Gamma \vdash x, y : \mathbf{El}_i(c)}{\Gamma \vdash \mathbf{eq}_{i,\Gamma}(c, x, y) : \mathbf{U}_i} \text{EQ/CODE (ETT)} \\
\\
\frac{\Gamma \vdash c : \mathbf{U}_i \quad \Gamma \vdash x, y : \mathbf{El}_i(c)}{\Gamma \vdash \mathbf{id}_{i,\Gamma}(c, x, y) : \mathbf{U}_i} \text{ID/CODE (ITT)} \quad \frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{unit}_{i,\Gamma} : \mathbf{U}_i} \text{UNIT/CODE} \\
\\
\frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{nat}_{i,\Gamma} : \mathbf{U}_i} \text{NAT/CODE} \quad \frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{void}_{i,\Gamma} : \mathbf{U}_i} \text{EMPTY/CODE} \quad \frac{\vdash \Gamma \text{ cx}}{\Gamma \vdash \mathbf{bool}_{i,\Gamma} : \mathbf{U}_i} \text{BOOL/CODE} \\
\\
\frac{\vdash \Gamma \text{ cx} \quad j < i}{\Gamma \vdash \mathbf{uni}_{\Gamma,i,j} : \mathbf{U}_i} \text{UNI/CODE} \quad \frac{\Gamma \vdash c : \mathbf{U}_i}{\Gamma \vdash \mathbf{lift}_{i,\Gamma}(c) : \mathbf{U}_{i+1}} \\
\\
\hline
\frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{U}_{\Gamma,i}[\gamma] = \mathbf{U}_{\Delta,i} \text{ type}} \quad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash a : \mathbf{U}_i}{\Delta \vdash \mathbf{El}_i(a)[\gamma] = \mathbf{El}_i(a[\gamma]) \text{ type}} \\
\\
\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash c_0 : \mathbf{U}_i \quad \Gamma.\mathbf{El}_i(c_0) \vdash c_1 : \mathbf{U}_i}{\Delta \vdash \mathbf{pi}(c_0, c_1)[\gamma] = \mathbf{pi}(c_0[\gamma], c_1[(\gamma \circ \mathbf{p}).\mathbf{q}]) : \mathbf{U}_i} \\
\\
\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash c_0 : \mathbf{U}_i \quad \Gamma.\mathbf{El}_i(c_0) \vdash c_1 : \mathbf{U}_i}{\Delta \vdash \mathbf{sig}(c_0, c_1)[\gamma] = \mathbf{sig}(c_0[\gamma], c_1[(\gamma \circ \mathbf{p}).\mathbf{q}]) : \mathbf{U}_i} \\
\\
\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash c : \mathbf{U}_i \quad \Gamma \vdash x, y : \mathbf{El}_i(c)}{\Delta \vdash \mathbf{eq}(c, x, y)[\gamma] = \mathbf{eq}(c[\gamma], x[\gamma], y[\gamma]) : \mathbf{U}_i} \text{(ETT)} \\
\\
\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash c : \mathbf{U}_i \quad \Gamma \vdash x, y : \mathbf{El}_i(c)}{\Delta \vdash \mathbf{id}(c, x, y)[\gamma] = \mathbf{id}(c[\gamma], x[\gamma], y[\gamma]) : \mathbf{U}_i} \text{(ITT)} \\
\\
\frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{unit}[\gamma] = \mathbf{unit} : \mathbf{U}_i} \quad \frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{nat}[\gamma] = \mathbf{nat} : \mathbf{U}_i} \quad \frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{void}[\gamma] = \mathbf{void} : \mathbf{U}_i} \\
\\
\frac{\Delta \vdash \gamma : \Gamma}{\Delta \vdash \mathbf{bool}[\gamma] = \mathbf{bool} : \mathbf{U}_i} \quad \frac{\Delta \vdash \gamma : \Gamma \quad j < i}{\Delta \vdash \mathbf{uni}_j[\gamma] = \mathbf{uni}_j : \mathbf{U}_i} \quad \frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash c : \mathbf{U}_i}{\Delta \vdash \mathbf{lift}(c)[\gamma] = \mathbf{lift}(c[\gamma]) : \mathbf{U}_{i+1}} \\
\\
\frac{\Gamma \vdash c_0 : \mathbf{U}_i \quad \Gamma.\mathbf{El}_i(c_0) \vdash c_1 : \mathbf{U}_i}{\Gamma \vdash \mathbf{El}_i(\mathbf{pi}(c_0, c_1)) = \mathbf{\Pi}(\mathbf{El}_i(c_0), \mathbf{El}_i(c_1)) \text{ type}}
\end{array}$$

$$\frac{\Gamma \vdash c_0 : \mathbf{U}_i \quad \Gamma.\mathbf{El}_i(c_0) \vdash c_1 : \mathbf{U}_i}{\Gamma \vdash \mathbf{El}_i(\mathbf{sig}(c_0, c_1)) = \Sigma(\mathbf{El}_i(c_0), \mathbf{El}_i(c_1)) \text{ type}}$$

$$\frac{\Gamma \vdash c : \mathbf{U}_i \quad \Gamma \vdash x, y : \mathbf{El}_i(c)}{\Gamma \vdash \mathbf{El}_i(\mathbf{eq}(c, x, y)) = \mathbf{Eq}(\mathbf{El}_i(c), x, y) \text{ type}} \text{ (ETT)}$$

$$\frac{\Gamma \vdash c : \mathbf{U}_i \quad \Gamma \vdash x, y : \mathbf{El}_i(c)}{\Gamma \vdash \mathbf{El}_i(\mathbf{id}(c, x, y)) = \mathbf{Id}(\mathbf{El}_i(c), x, y) \text{ type}} \text{ (ITT)} \quad \frac{}{\Gamma \vdash \mathbf{El}_i(\mathbf{unit}) = \mathbf{Unit} \text{ type}}$$

$$\frac{}{\Gamma \vdash \mathbf{El}_i(\mathbf{nat}) = \mathbf{Nat} \text{ type}} \quad \frac{}{\Gamma \vdash \mathbf{El}_i(\mathbf{void}) = \mathbf{Void} \text{ type}} \quad \frac{}{\Gamma \vdash \mathbf{El}_i(\mathbf{bool}) = \mathbf{Bool} \text{ type}}$$

$$\frac{j < i}{\Gamma \vdash \mathbf{El}_i(\mathbf{uni}_j) = \mathbf{U}_j \text{ type}}$$

$$\frac{\Gamma \vdash c : \mathbf{U}_i}{\Gamma \vdash \mathbf{El}_{i+1}(\mathbf{lift}(c)) = \mathbf{El}_i(c) \text{ type}}$$

Solutions to selected exercises

Solution 2.2. Any substitution γ into $\Gamma.A$ is of the form $(\mathbf{p} \circ \gamma).\mathbf{q}[\gamma]$, which by our hypothesis is equal to $\mathbf{id}.\mathbf{q}[\gamma]$. We can apply this substitution to a variable, obtaining the term $\Gamma \vdash \mathbf{q}[\mathbf{id}.\mathbf{q}[\gamma]] = \mathbf{q}[\gamma] : A[\mathbf{id}]$ as required. Conversely, any term $\Gamma \vdash a : A$ determines a substitution $\Gamma \vdash \mathbf{id}.a : \Gamma.A$ that satisfies $\mathbf{p} \circ (\mathbf{id}.a) = \mathbf{id}$. One round-trip follows from the previously noted equation, and the other from $\mathbf{q}[\mathbf{id}.a] = a$.

Solution 2.3. To show

$$\frac{\exists \delta : \Delta \quad \Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type} \quad \Delta \vdash a : A[\gamma]}{\Delta \vdash (\gamma.a) \circ \delta = (\gamma \circ \delta).a[\delta] : \Gamma.A} \Rightarrow$$

we calculate as follows:

$$\begin{aligned} (\gamma.a) \circ \delta &= (\mathbf{p} \circ (\gamma.a) \circ \delta).(\mathbf{q}[(\gamma.a) \circ \delta]) \\ &= (\gamma \circ \delta).(\mathbf{q}[\gamma.a][\delta]) \\ &= (\gamma \circ \delta).(a[\delta]) \end{aligned}$$

Solution 2.4. We define $\gamma.A := (\gamma \circ \mathbf{p}).\mathbf{q}$, i.e., the extension of the substitution $\Delta.A[\gamma] \vdash \gamma \circ \mathbf{p} : \Gamma$ by the variable $\Delta.A[\gamma] \vdash \mathbf{q} : A[\gamma \circ \mathbf{p}]$.

Solution 2.5. In the forward direction, we send $\Delta \vdash \gamma : \Gamma.A$ to the pair of $\mathbf{p} \circ \gamma$ and $\mathbf{q}[\gamma]$; in the reverse direction, we send pairs of γ_0 and a to the substitution $\gamma_0.a$. One round-trip follows from $\gamma = (\mathbf{p} \circ \gamma).\mathbf{q}[\gamma]$ and the other from $\mathbf{p} \circ (\gamma_0.a) = \gamma_0$ and $\mathbf{q}[\gamma_0.a] = a$.

Solution 2.8. Below are the formation, introduction, and elimination rules for non-dependent functions, along with their definitions in terms of Π -types:

$$\frac{\Gamma \vdash A \text{ type} \quad \Gamma \vdash B \text{ type}}{\Gamma \vdash A \rightarrow B := \Pi(A, B[\mathbf{p}]) \text{ type}} \quad \frac{\Gamma \vdash A \text{ type} \quad \Gamma.A \vdash b : B[\mathbf{p}]}{\Gamma \vdash \lambda \mathbf{q}.b := \lambda(b) : A \rightarrow B}$$

$$\frac{\Gamma \vdash a : A \quad \Gamma \vdash B \text{ type} \quad \Gamma \vdash f : A \rightarrow B}{\Gamma \vdash f a := \mathbf{app}(f, a) : B}$$

Note that B must be weakened, and the elimination rule is meta-well-typed because $B[\mathbf{p} \circ (\mathbf{id}.a)] = B$. The β - and η -rules are immediate.

Solution 2.16. The only non-trivial presupposition to check is $\Delta \vdash \mathbf{pair}(a[\gamma], b[\gamma]) : \Sigma(A, B)[\gamma]$. By the substitution rule for Σ , we have $\Sigma(A, B)[\gamma] = \Sigma(A[\gamma], B[\gamma.A])$. The first component of the \mathbf{pair} is thus well-typed by $\Delta \vdash a[\gamma] : A[\gamma]$. For the second component, we must show $\Delta \vdash b[\gamma] : B[\gamma.A][\mathbf{id}.a[\gamma]]$. By applying γ to the typing premise for b we obtain $\Delta \vdash b[\gamma] : B[\mathbf{id}.a][\gamma]$, so it suffices to show $(\gamma.A) \circ (\mathbf{id}.a[\gamma]) = (\mathbf{id}.a) \circ \gamma$:

$$\begin{aligned}
& (\gamma.A) \circ (\mathbf{id}.a[\gamma]) \\
&= ((\gamma \circ \mathbf{p}).\mathbf{q}) \circ (\mathbf{id}.a[\gamma]) && \text{by Exercise 2.4} \\
&= (\gamma \circ \mathbf{p} \circ (\mathbf{id}.a[\gamma])).\mathbf{q}[\mathbf{id}.a[\gamma]] && \text{by Exercise 2.3} \\
&= (\gamma \circ \mathbf{id}).a[\gamma] \\
&= (\mathbf{id} \circ \gamma).a[\gamma] \\
&= (\mathbf{id}.a) \circ \gamma && \text{by Exercise 2.3}
\end{aligned}$$

Solution 2.17. The substitution rule is somewhat odd:

$$\frac{\Delta \vdash \gamma : \Gamma \quad \Gamma \vdash A \text{ type} \quad \Gamma.A \vdash B \text{ type} \quad \Gamma \vdash f : \Pi(A, B)}{\Delta.A[\gamma] \vdash \lambda^{-1}(f[\gamma]) = \lambda^{-1}(f)[\gamma.A] : B[\gamma.A]}$$

We prove it as follows:

$$\begin{aligned}
& \lambda^{-1}(f[\gamma]) \\
&= \lambda^{-1}(\lambda(\lambda^{-1}(f))[\gamma]) && \text{by } f = \lambda(\lambda^{-1}(f)) \\
&= \lambda^{-1}(\lambda(\lambda^{-1}(f)[\gamma.A])) && \text{by substitution for } \lambda \\
&= \lambda^{-1}(f)[\gamma.A] && \text{by } \lambda^{-1}(\lambda(\dots)) = \dots
\end{aligned}$$

Solution 2.20. The elimination principle corresponds to the forward map $\iota_\Gamma : \mathbf{Tm}(\Gamma, \mathbf{Unit}) \rightarrow \{\star\}$. This tells us that from $\Gamma \vdash a : \mathbf{Unit}$ we can obtain an element of $\{\star\}$, a principle which contains no useful information. The substitution rule for \mathbf{tt} states that $\Delta \vdash \mathbf{tt}[\gamma] = \mathbf{tt} : \mathbf{Unit}$, but this follows already from the η principle. Equivalently, in terms of the natural isomorphism, the forward maps ι_Γ are natural “for free” because *all* elements of $\{\star\}$ are equal; thus the backward maps ι_Γ^{-1} (which determine \mathbf{tt}) are also automatically natural.

Solution 3.1.

$$\frac{\Gamma \vdash \tau_0 \text{ type} \rightsquigarrow A \quad \Gamma.A \vdash \tau_1 \text{ type} \rightsquigarrow B \quad \Gamma \vdash e_0 : A \rightsquigarrow a \quad \Gamma \vdash e_1 : B[\mathbf{id}.a] \rightsquigarrow b \quad \Gamma \vdash C = \Sigma(A, B) \text{ type}}{\Gamma \vdash (\mathbf{pair} \tau_0 \tau_1 e_0 e_1) : C \rightsquigarrow \mathbf{pair}_{\Gamma, A, B}(a, b)}$$

Solution 3.7. By Slogan 3.2.7, we check $(\text{pair } e_0 e_1)$ and synthesize $(\text{fst } e)$ and $(\text{snd } e)$.

$$\frac{\text{unSigma}(C) = (A, B) \quad \Gamma \vdash e_0 \Leftarrow A \rightsquigarrow a \quad \Gamma \vdash e_1 \Leftarrow B[\text{id}.a] \rightsquigarrow b}{\Gamma \vdash (\text{pair } e_0 e_1) \Leftarrow C \rightsquigarrow \text{pair}(a, b)}$$

$$\frac{\Gamma \vdash e \Rightarrow C \rightsquigarrow p \quad \text{unSigma}(C) = (A, B)}{\Gamma \vdash (\text{fst } e) \Rightarrow A \rightsquigarrow \text{fst}(p)} \quad \frac{\Gamma \vdash e \Rightarrow C \rightsquigarrow p \quad \text{unSigma}(C) = (A, B)}{\Gamma \vdash (\text{snd } e) \Rightarrow B[\text{id}.\text{fst}(p)] \rightsquigarrow \text{snd}(p)}$$

In the above rules, unSigma is an algorithm that inverts Σ -types: given $\Gamma \vdash C$ type it returns the unique pair of types A, B for which $\Gamma \vdash C = \Sigma(A, B)$ type, if they exist.

Solution 3.8. The fixed-point of the identity function $\mathbf{Void} \rightarrow \mathbf{Void}$ is a closed proof of \mathbf{Void} :

$$\frac{1 \vdash \mathbf{Void} \text{ type} \quad 1.\mathbf{Void} \vdash q : \mathbf{Void}}{1 \vdash \text{fix}(q) : \mathbf{Void}}$$

Solution 3.9. Suppose there is a model \mathcal{M} for which $\text{Tm}_{\mathcal{M}}(\mathbf{1}_{\mathcal{M}}, \mathbf{Bool}_{\mathcal{M}})$ has exactly two elements. By Theorem 3.4.5 there is a function $\text{Tm}_f(\mathbf{1}, \mathbf{Bool}) : \text{Tm}(\mathbf{1}, \mathbf{Bool}) \rightarrow \text{Tm}_{\mathcal{M}}(\mathbf{1}_{\mathcal{M}}, \mathbf{Bool}_{\mathcal{M}})$, but this does not allow us to conclude that $\text{Tm}(\mathbf{1}, \mathbf{Bool})$ has exactly two elements!

In Theorem 3.4.7, the existence of a function $X \rightarrow \emptyset$ allowed us to observe that $X = \emptyset$, but the existence of a function $X \rightarrow \{\star, \star'\}$ does not imply X has exactly two elements.

Solution 4.7. Define $c_a = c a$.

Solution 4.8. Define $q = \text{uniq}(a_1, p)$.

Solution 4.9. Define $c_b = \text{subst } C_a q c_a$.

Solution 4.10. We have $c_b : C_a(b, p)$ but $C_a(b, p) = C a b p$ by definition. We define j as follows:

$$j : \{A : \mathbf{U}\} (C : (a b : A) \rightarrow \mathbf{Id}(A, a, b) \rightarrow \mathbf{U}) \rightarrow ((a : A) \rightarrow C a a \text{ refl}) \rightarrow \\ (a b : A) (p : \mathbf{Id}(A, a, b)) \rightarrow C a b p \\ j \{A\} C c a b p = \text{subst } (\lambda x \rightarrow C a (\text{fst } x) (\text{snd } x)) (\text{uniq}(b, p)) (c a)$$

Solution 4.11.

$$\begin{aligned}
& \text{j } C \text{ c } a \text{ a refl} \\
& = \text{subst } (\lambda x \rightarrow C \text{ a } (\text{fst } x) (\text{snd } x)) (\text{uniq } (a, \text{refl})) (c \text{ a}) && \text{by Exercise 4.10} \\
& = \text{subst } (\lambda x \rightarrow C \text{ a } (\text{fst } x) (\text{snd } x)) \text{refl } (c \text{ a}) && \text{by uniq def.eq.} \\
& = c \text{ a} && \text{by subst def.eq.}
\end{aligned}$$

Solution 5.1. Fix $h : (a : a)(b_0, b_1 : B a) \rightarrow \text{Id}(b a, b_0, b_1)$ along with $f_0, f_1 : (a : a) \rightarrow B a$. The identification between them is given by $\text{funext}(\lambda a \rightarrow h a (f_0 a) (f_1 a))$.

Bibliography

- [Abe13] Andreas Abel. “Normalization by Evaluation: Dependent Types and Impredicativity”. Habilitation thesis. Ludwig-Maximilians-Universität München, 2013. URL: <http://www2.tcs.ifi.lmu.de/~abel/habil.pdf>.
- [ACD07] Andreas Abel, Thierry Coquand, and Peter Dybjer. “Normalization by Evaluation for Martin-Löf Type Theory with Typed Equality Judgements”. In: *22nd Annual IEEE Symposium on Logic in Computer Science (LICS 2007)*. July 2007, pp. 3–12. DOI: [10.1109/LICS.2007.33](https://doi.org/10.1109/LICS.2007.33).
- [AFH17] Carlo Angiuli, Kuen-Bang Hou (Favonia), and Robert Harper. *Computational Higher Type Theory III: Univalent Universes and Exact Equality*. Preprint. Dec. 2017. arXiv: [1712.01800 \[cs.LO\]](https://arxiv.org/abs/1712.01800).
- [AFH18] Carlo Angiuli, Kuen-Bang Hou (Favonia), and Robert Harper. “Cartesian Cubical Computational Type Theory: Constructive Reasoning with Paths and Equalities”. In: *27th EACSL Annual Conference on Computer Science Logic (CSL 2018)*. Ed. by Dan Ghica and Achim Jung. Vol. 119. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018, 6:1–6:17. ISBN: 978-3-95977-088-0. DOI: [10.4230/LIPIcs.CSL.2018.6](https://doi.org/10.4230/LIPIcs.CSL.2018.6).
- [Agda] The Agda Development Team. *The Agda Programming Language*. 2020. URL: <http://wiki.portal.chalmers.se/agda/pmwiki.php>.
- [Ahr+22] Benedikt Ahrens, Paige Randall North, Michael Shulman, and Dimitris Tsementzis. *The Univalence Principle*. 2022. arXiv: [2102.06275 \[math.CT\]](https://arxiv.org/abs/2102.06275).
- [AJ21] Mathieu Anel and André Joyal. “Topologie”. In: *New Spaces in Mathematics: Formal and Conceptual Reflections*. Ed. by Mathieu Anel and Gabriel Editors Catren. Cambridge University Press, 2021, pp. 155–257.
- [AK16] Thorsten Altenkirch and Ambrus Kaposi. “Type theory in type theory using quotient inductive types”. In: *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL 2016. New York, NY, USA: ACM, 2016, pp. 18–29. ISBN: 9781450335492. DOI: [10.1145/2837614.2837638](https://doi.org/10.1145/2837614.2837638).
- [Alt+01] T. Altenkirch, P. Dybjer, M. Hofmann, and P. Scott. “Normalization by evaluation for typed lambda calculus with coproducts”. In: *Proceedings of the 16th Annual IEEE Symposium on Logic in Computer Science (LICS 2001)*. 2001, pp. 303–310. DOI: [10.1109/LICS.2001.932506](https://doi.org/10.1109/LICS.2001.932506).

- [Alt23] Thorsten Altenkirch. “Should Type Theory Replace Set Theory as the Foundation of Mathematics?” In: *Global Philosophy* 33.21 (2023). DOI: [10.1007/s10516-023-09676-0](https://doi.org/10.1007/s10516-023-09676-0).
- [AMB13] Guillaume Allais, Conor McBride, and Pierre Boutillier. “New equations for neutral terms: a sound and complete decision procedure, formalized”. In: *Proceedings of the 2013 ACM SIGPLAN Workshop on Dependently-Typed Programming*. DTP ’13. New York, NY, USA: ACM, 2013, pp. 13–24. ISBN: 9781450323840. DOI: [10.1145/2502409.2502411](https://doi.org/10.1145/2502409.2502411).
- [AMS07] Thorsten Altenkirch, Conor McBride, and Wouter Swierstra. “Observational Equality, Now!” In: *Proceedings of the 2007 Workshop on Programming Languages Meets Program Verification*. PLPV ’07. New York, NY, USA: ACM, 2007, pp. 57–68. ISBN: 978-1-59593-677-6. DOI: [10.1145/1292597.1292608](https://doi.org/10.1145/1292597.1292608).
- [Ane19] Mathieu Anel. “Descent & Univalence”. HoTTEST talk. 2019. URL: <http://mathieu.anel.free.fr/mat/doc/Anel-2019-HoTTEST.pdf>.
- [Ang19] Carlo Angiuli. “Computational Semantics of Cartesian Cubical Type Theory”. PhD thesis. Carnegie Mellon University, Sept. 2019. URL: <http://reports-archive.adm.cs.cmu.edu/anon/2019/CMU-CS-19-127.pdf>.
- [Ang+21] Carlo Angiuli, Guillaume Brunerie, Thierry Coquand, Robert Harper, Kuen-Bang Hou (Favonia), and Daniel R. Licata. “Syntax and models of Cartesian cubical type theory”. In: *Mathematical Structures in Computer Science* 31.4 (2021). Special issue on Homotopy Type Theory and Univalent Foundations, pp. 424–468. DOI: [10.1017/S0960129521000347](https://doi.org/10.1017/S0960129521000347).
- [AÖV17] Andreas Abel, Joakim Öhman, and Andrea Vezzosi. “Decidability of conversion for type theory in type theory”. In: *Proceedings of the ACM on Programming Languages* 2.POPL (Dec. 2017), 23:1–23:29. DOI: [10.1145/3158111](https://doi.org/10.1145/3158111).
- [Asp95] David Aspinall. “Subtyping with singleton types”. In: *Computer Science Logic*. Ed. by Leszek Pacholski and Jerzy Tiuryn. Berlin, Heidelberg: Springer Berlin Heidelberg, 1995, pp. 1–15. ISBN: 978-3-540-49404-1.
- [Aug99] Lennart Augustsson. “Cayenne — A Language with Dependent Types”. In: *Advanced Functional Programming*. Ed. by S. Doaitse Swierstra, José N. Oliveira, and Pedro R. Henriques. Berlin, Heidelberg: Springer Berlin Heidelberg, 1999, pp. 240–267. ISBN: 978-3-540-48506-3. DOI: [10.1007/10704973_6](https://doi.org/10.1007/10704973_6).
- [AW09] Steve Awodey and Michael A. Warren. “Homotopy theoretic models of identity types”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 146.1 (Jan. 2009), pp. 45–55. ISSN: 0305-0041. DOI: [10.1017/S0305004108001783](https://doi.org/10.1017/S0305004108001783).

- [Awo18] Steve Awodey. “Natural models of homotopy type theory”. In: *Mathematical Structures in Computer Science* 28.2 (2018), pp. 241–286. DOI: [10.1017/S0960129516000268](https://doi.org/10.1017/S0960129516000268).
- [Bar91] Henk Barendregt. “Introduction to generalized type systems”. In: *Journal of Functional Programming* 1.2 (1991), pp. 125–154. DOI: [10.1017/S095679680020025](https://doi.org/10.1017/S095679680020025).
- [Bau+21] Andrej Bauer, Gaëtan Gilbert, Philipp G. Haselwarter, Anja Petković, Matija Pretnar, and Chris Stone. *Andromeda: Your type theory à la Martin-Löf*. 2021. URL: <https://www.andromeda-prover.org/>.
- [Bra13] Edwin Brady. “Idris, a general-purpose dependently typed programming language: Design and implementation”. In: *Journal of Functional Programming* 23.5 (2013), pp. 552–593. DOI: [10.1017/S095679681300018X](https://doi.org/10.1017/S095679681300018X).
- [Bra17] Edwin Brady. *Type-Driven Development with Idris*. Manning Publications, 2017. ISBN: 9781617293023.
- [Bru16] Guillaume Brunerie. “On the homotopy groups of spheres in homotopy type theory”. PhD thesis. Université Nice Sophia Antipolis, 2016. URL: <http://arxiv.org/abs/1606.05916>.
- [Bru18] Guillaume Brunerie. “The James Construction and $\pi_4(\mathbb{S}^3)$ in Homotopy Type Theory”. In: *Journal of Automated Reasoning* 63.2 (June 2018), 255–284. ISSN: 1573-0670. DOI: [10.1007/s10817-018-9468-2](https://doi.org/10.1007/s10817-018-9468-2). URL: <http://dx.doi.org/10.1007/s10817-018-9468-2>.
- [Bru72] N. G. de Bruijn. “Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem”. In: *Indagationes Mathematicae* 75.5 (1972), pp. 381–392. ISSN: 1385-7258. DOI: [10.1016/1385-7258\(72\)90034-0](https://doi.org/10.1016/1385-7258(72)90034-0).
- [Car86] John Cartmell. “Generalised algebraic theories and contextual categories”. In: *Annals of Pure and Applied Logic* 32 (1986), pp. 209–243. ISSN: 0168-0072. DOI: [10.1016/0168-0072\(86\)90053-9](https://doi.org/10.1016/0168-0072(86)90053-9).
- [CCD17] Simon Castellan, Pierre Clairambault, and Peter Dybjer. “Undecidability of Equality in the Free Locally Cartesian Closed Category”. In: *Logical Methods in Computer Science* Volume 13, Issue 4 (Nov. 2017). DOI: [10.23638/LMCS-13\(4:22\)2017](https://doi.org/10.23638/LMCS-13(4:22)2017).
- [CCD21] Simon Castellan, Pierre Clairambault, and Peter Dybjer. “Categories with Families: Untyped, Simply Typed, and Dependently Typed”. In: *Joachim Lambek: The Interplay of Mathematics, Logic, and Linguistics*. Ed. by Claudia Casadio and Philip J. Scott. Cham: Springer International Publishing, 2021, pp. 135–180. ISBN: 978-3-030-66545-6. DOI: [10.1007/978-3-030-66545-6_5](https://doi.org/10.1007/978-3-030-66545-6_5).

- [CCHM18] Cyril Cohen, Thierry Coquand, Simon Huber, and Anders Mörtberg. “Cubical Type Theory: A Constructive Interpretation of the Univalence Axiom”. In: *21st International Conference on Types for Proofs and Programs (TYPES 2015)*. Ed. by Tarmo Uustalu. Vol. 69. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018, 5:1–5:34. ISBN: 978-3-95977-030-9. DOI: [10.4230/LIPIcs.TYPES.2015.5](https://doi.org/10.4230/LIPIcs.TYPES.2015.5).
- [CDP14] Jesper Cockx, Dominique Devriese, and Frank Piessens. “Pattern Matching without K”. In: *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming*. ICFP ’14. New York, NY, USA: ACM, 2014, pp. 257–268. ISBN: 978-1-4503-2873-9. DOI: [10.1145/2628136.2628139](https://doi.org/10.1145/2628136.2628139).
- [CH19] Evan Cavallo and Robert Harper. “Higher Inductive Types in Cubical Computational Type Theory”. In: *Proceedings of the ACM on Programming Languages* 3.POPL (Jan. 2019), 1:1–1:27. ISSN: 2475-1421. DOI: [10.1145/3290314](https://doi.org/10.1145/3290314).
- [CH88] Thierry Coquand and Gérard Huet. “The Calculus of Constructions”. In: *Information and Computation* 76.2 (1988), pp. 95–120. ISSN: 0890-5401. DOI: [10.1016/0890-5401\(88\)90005-3](https://doi.org/10.1016/0890-5401(88)90005-3).
- [Chr19] David Christiansen. *Checking Dependent Types with Normalization by Evaluation: A Tutorial*. Online. 2019. URL: <https://davidchristiansen.dk/tutorials/nbe/>.
- [Chr23] David Thrane Christiansen. *Functional Programming in Lean*. 2023. URL: https://lean-lang.org/functional_programming_in_lean/.
- [Con+85] R. L. Constable, S. F. Allen, H. M. Bromley, W. R. Cleaveland, J. F. Cremer, R. W. Harper, D. J. Howe, T. B. Knoblock, N. P. Mendler, P. Panangaden, J. T. Sasaki, and S. F. Smith. *Implementing Mathematics with the Nuprl Proof Development Environment*. Prentice-Hall, 1985. URL: <http://www.nuprl.org/book/>.
- [Coq] The Coq Development Team. *The Coq Proof Assistant*. 2020. URL: <https://www.coq.inria.fr>.
- [Coq+09] Thierry Coquand, Yoshiki Kinoshita, Bengt Nordström, and Makoto Takeyama. “A simple type-theoretic language: Mini-TT”. In: *From Semantics to Computer Science: Essays in Honour of Gilles Kahn*. Ed. by Yves Bertot, Gérard Huet, Jean-Jacques Lévy, and Gordon Plotkin. Cambridge University Press, 2009, pp. 139–164.
- [Coq13] Thierry Coquand. “Presheaf model of type theory”. Unpublished note. 2013. URL: <http://www.cse.chalmers.se/~coquand/presheaf.pdf>.
- [Coq14] Thierry Coquand. *A remark on singleton types*. Online. Mar. 2014. URL: <https://www.cse.chalmers.se/~coquand/singl.pdf>.

- [Coq19] Thierry Coquand. “Canonicity and normalization for dependent type theory”. In: *Theoretical Computer Science* 777 (2019). In memory of Maurice Nivat, a founding father of Theoretical Computer Science - Part I, pp. 184–191. ISSN: 0304-3975. DOI: [10.1016/j.tcs.2019.01.015](https://doi.org/10.1016/j.tcs.2019.01.015).
- [Coq86] Thierry Coquand. “An Analysis of Girard’s Paradox”. In: *Proceedings of the First Annual IEEE Symposium on Logic in Computer Science (LICS 1986)*. Cambridge, MA, USA: IEEE Computer Society Press, 1986, pp. 227–236.
- [Coq91] Thierry Coquand. “An algorithm for testing conversion in type theory”. In: *Logical Frameworks*. USA: Cambridge University Press, 1991, pp. 255–279. ISBN: 0521413001.
- [Coq92] Thierry Coquand. “Pattern matching with dependent types”. In: *Types for proofs and programs*. Ed. by Bengt Nordström, Kent Petersson, and Gordon Plotkin. 1992.
- [Coq96] Thierry Coquand. “An algorithm for type-checking dependent types”. In: *Science of Computer Programming* 26.1 (1996), pp. 167–177. ISSN: 0167-6423. DOI: [https://doi.org/10.1016/0167-6423\(95\)00021-6](https://doi.org/10.1016/0167-6423(95)00021-6).
- [CP90a] Thierry Coquand and Christine Paulin. “Inductively defined types”. In: *COLOG-88*. Ed. by Per Martin-Löf and Grigori Mints. Berlin, Heidelberg: Springer Berlin Heidelberg, 1990, pp. 50–66. ISBN: 978-3-540-46963-6. DOI: [10.1007/3-540-52335-9_47](https://doi.org/10.1007/3-540-52335-9_47).
- [CP90b] Thierry Coquand and Christine Paulin-Mohring. “Inductively defined types”. In: *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 1990, pp. 50–66. ISBN: 9783540469636. DOI: [10.1007/3-540-52335-9_47](https://doi.org/10.1007/3-540-52335-9_47).
- [Cro94] Roy L. Crole. *Categories for Types*. Cambridge: Cambridge University Press, 1994. DOI: [10.1017/CB09781139172707](https://doi.org/10.1017/CB09781139172707).
- [Dij17] Gabe Dijkstra. “Quotient inductive-inductive definitions”. PhD thesis. University of Nottingham, 2017. URL: <https://eprints.nottingham.ac.uk/42317/>.
- [Dow93] Gilles Dowek. “The undecidability of typability in the Lambda-Pi-calculus”. In: *Typed Lambda Calculi and Applications*. Ed. by Marc Bezem and Jan Friso Groote. Berlin, Heidelberg: Springer Berlin Heidelberg, 1993, pp. 139–145. ISBN: 978-3-540-47586-6.
- [Dyb00] Peter Dybjer. “A General Formulation of Simultaneous Inductive-Recursive Definitions in Type Theory”. In: *The Journal of Symbolic Logic* 65.2 (2000), pp. 525–549. ISSN: 00224812. DOI: [10.2307/2586554](https://doi.org/10.2307/2586554).
- [Dyb94] Peter Dybjer. “Inductive families”. In: *Formal Aspects of Computing* 6.4 (July 1994), pp. 440–465. ISSN: 0934-5043. DOI: [10.1007/BF01211308](https://doi.org/10.1007/BF01211308).

- [Dyb96] Peter Dybjer. “Internal type theory”. In: *Types for Proofs and Programs (TYPES 1995)*. Ed. by Stefano Berardi and Mario Coppo. Vol. 1158. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg, 1996, pp. 120–134. ISBN: 978-3-540-70722-6. DOI: [10.1007/3-540-61780-9_66](https://doi.org/10.1007/3-540-61780-9_66).
- [Esc14] Martín Escardó. *Comment on “Generalize 7.2.2 and simplify encode-decode”*. Online. Comment on a GitHub issue. 2014. URL: <https://github.com/HoTT/book/issues/718#issuecomment-65378867>.
- [Esc+14] Martín H. Escardó and contributors. *TypeTopology*. Agda development. URL: <https://github.com/martinescardo/TypeTopology>.
- [FAM23] Kuen-Bang Hou (Favonia), Carlo Angiuli, and Reed Mullanix. “An Order-Theoretic Analysis of Universe Polymorphism”. In: *Proceedings of the ACM on Programming Languages* 7.POPL (Jan. 2023). DOI: [10.1145/3571250](https://doi.org/10.1145/3571250).
- [Fav+16] Kuen-Bang Hou (Favonia), Eric Finster, Daniel R. Licata, and Peter LeFanu Lumsdaine. “A Mechanization of the Blakers-Massey Connectivity Theorem in Homotopy Type Theory”. In: *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science*. LICS ’16. New York, NY, USA: ACM, 2016, pp. 565–574. ISBN: 978-1-4503-4391-6. DOI: [10.1145/2933575.2934545](https://doi.org/10.1145/2933575.2934545).
- [FC18] Daniel P. Friedman and David Thrane Christiansen. *The Little Typer*. The MIT Press, 2018. ISBN: 9780262536431.
- [Fio02] Marcelo Fiore. “Semantic Analysis of Normalisation by Evaluation for Typed Lambda Calculus”. In: *Proceedings of the 4th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*. PPDP ’02. Pittsburgh, PA, USA: ACM, 2002, pp. 26–37. ISBN: 1-58113-528-9. DOI: [10.1145/571157.571161](https://doi.org/10.1145/571157.571161).
- [Fre78] Peter Freyd. “On proving that 1 is an indecomposable projective in various free categories”. Unpublished note. 1978.
- [Gil+19] Gaëtan Gilbert, Jesper Cockx, Matthieu Sozeau, and Nicolas Tabareau. “Definitional proof-irrelevance without K”. In: *Proceedings of the ACM on Programming Languages* 3.POPL (Jan. 2019). DOI: [10.1145/3290316](https://doi.org/10.1145/3290316).
- [Gir99] Jean-Yves Girard. “On the Meaning of Logical Rules I: Syntax Versus Semantics”. In: *Computational Logic*. Springer Berlin Heidelberg, 1999, pp. 215–272. ISBN: 9783642586224. DOI: [10.1007/978-3-642-58622-4_7](https://doi.org/10.1007/978-3-642-58622-4_7). URL: http://dx.doi.org/10.1007/978-3-642-58622-4_7.
- [GK16] David Gepner and Joachim Kock. “Univalence in locally cartesian closed ∞ -categories”. In: *Forum Mathematicum* 29.3 (Nov. 2016), pp. 617–652. ISSN: 0933-7741. DOI: [10.1515/forum-2015-0228](https://doi.org/10.1515/forum-2015-0228).

- [Gle14] Tamara von Glehn. “Polynomials and Models of Type Theory”. PhD thesis. University of Cambridge, 2014. DOI: [10.17863/CAM.16245](https://doi.org/10.17863/CAM.16245).
- [GLT89] Jean-Yves Girard, Yves Lafont, and Paul Taylor. *Proofs and Types*. Cambridge Tracts in Theoretical Computer Science 7. Cambridge University Press, 1989.
- [GMW79] Michael J. Gordon, Arthur J. Milner, and Christopher P. Wadsworth. *Edinburgh LCF*. Springer Berlin Heidelberg, 1979. ISBN: 9783540385264. DOI: [10.1007/3-540-09724-4](https://doi.org/10.1007/3-540-09724-4).
- [Gra09] Johan Georg Granström. “Reference and Computation in Intuitionistic Type Theory”. PhD thesis. Uppsala University, 2009. URL: https://intuitionistic.files.wordpress.com/2010/07/theses_published_uppsala.pdf.
- [Gra+21] Daniel Gratzer, G. A. Kavvos, Andreas Nuyts, and Lars Birkedal. “Multimodal Dependent Type Theory”. In: *Logical Methods in Computer Science* 17.3 (July 2021). DOI: [10.46298/lmcs-17\(3:11\)2021](https://doi.org/10.46298/lmcs-17(3:11)2021).
- [Gra+22] Daniel Gratzer, Evan Cavallo, G. A. Kavvos, Adrien Guatto, and Lars Birkedal. “Modalities and Parametric Adjoints”. In: *ACM Transactions on Computational Logic* 23.3 (Apr. 2022). ISSN: 1529-3785. DOI: [10.1145/3514241](https://doi.org/10.1145/3514241).
- [Gra23] Daniel Gratzer. “Syntax and semantics of modal type theory”. PhD thesis. Aarhus University, 2023. URL: [https://pure.au.dk/portal/en/publications/syntax-and-semantics-of-modal-type-theory\(694f77d2-47d3-4986-bb82-129b8d96206e\).html](https://pure.au.dk/portal/en/publications/syntax-and-semantics-of-modal-type-theory(694f77d2-47d3-4986-bb82-129b8d96206e).html).
- [GTL89] Jean-Yves Girard, Paul Taylor, and Yves Lafont. *Proofs and types*. USA: Cambridge University Press, 1989. ISBN: 0521371813.
- [Har09] John Harrison. “HOL Light: An Overview”. In: *Theorem Proving in Higher Order Logics*. Ed. by Stefan Berghofer, Tobias Nipkow, Christian Urban, and Makarius Wenzel. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 60–66. ISBN: 978-3-642-03359-9. DOI: [10.1007/978-3-642-03359-9_4](https://doi.org/10.1007/978-3-642-03359-9_4).
- [Har16] Robert Harper. *Practical Foundations for Programming Languages*. Second Edition. Cambridge University Press, 2016. ISBN: 9781107150300. DOI: [10.1017/CB09781316576892](https://doi.org/10.1017/CB09781316576892).
- [Has21] Philipp Haselwarter. “Effective Metatheory for Type Theory”. PhD thesis. University of Ljubljana, 2021. URL: <https://repozitorij.uni-lj.si/IzpisGradiva.php?id=134439>.
- [HHP93] Robert Harper, Furio Honsell, and Gordon Plotkin. “A framework for defining logics”. In: *J. ACM* 40.1 (Jan. 1993), pp. 143–184. ISSN: 0004-5411. DOI: [10.1145/138027.138060](https://doi.org/10.1145/138027.138060).

- [HM95] Robert Harper and Greg Morrisett. “Compiling Polymorphism Using Intensional Type Analysis”. In: *Proceedings of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL ’95. New York, NY, USA: ACM, 1995, pp. 130–141. ISBN: 0897916921. DOI: [10.1145/199448.199475](https://doi.org/10.1145/199448.199475).
- [Hof95a] Martin Hofmann. “Extensional concepts in intensional type theory”. PhD thesis. University of Edinburgh, July 1995. URL: <http://www.lfcs.inf.ed.ac.uk/reports/95/ECS-LFCS-95-327/>.
- [Hof95b] Martin Hofmann. “On the interpretation of type theory in locally cartesian closed categories”. In: *8th Workshop, Computer Science Logic (CSL 1994)*. Ed. by Leszek Pacholski and Jerzy Tiuryn. Vol. 933. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg, 1995, pp. 427–441. ISBN: 978-3-540-49404-1. DOI: [10.1007/BFb0022273](https://doi.org/10.1007/BFb0022273).
- [Hof97] Martin Hofmann. “Syntax and Semantics of Dependent Types”. In: *Semantics and Logics of Computation*. Ed. by Andrew M. Pitts and P. Dybjer. Publications of the Newton Institute. Cambridge University Press, 1997, pp. 79–130. DOI: [10.1017/CB09780511526619.004](https://doi.org/10.1017/CB09780511526619.004).
- [HS98] Martin Hofmann and Thomas Streicher. “The groupoid interpretation of type theory”. In: *Twenty-Five Years of Constructive Type Theory*. Ed. by Giovanni Sambin and Jan Smith. Vol. 36. Oxford Logic Guides. Oxford University Press, 1998, pp. 83–111.
- [Hub18] Simon Huber. “Canonicity for Cubical Type Theory”. In: *Journal of Automated Reasoning* (June 2018). ISSN: 1573-0670. DOI: [10.1007/s10817-018-9469-1](https://doi.org/10.1007/s10817-018-9469-1).
- [Hur95] Antonius J. C. Hurkens. “A simplification of Girard’s paradox”. In: *Typed Lambda Calculi and Applications*. Ed. by Mariangiola Dezani-Ciancaglini and Gordon Plotkin. TLCA 1995. Berlin, Heidelberg: Springer Berlin Heidelberg, 1995, pp. 266–278. ISBN: 978-3-540-49178-1. DOI: [10.1007/BFb0014058](https://doi.org/10.1007/BFb0014058).
- [Hyl82] J.M.E. Hyland. “The Effective Topos”. In: *Studies in Logic and the Foundations of Mathematics*. Elsevier, 1982, pp. 165–216. DOI: [10.1016/s0049-237x\(09\)70129-6](https://doi.org/10.1016/s0049-237x(09)70129-6). URL: [http://dx.doi.org/10.1016/S0049-237X\(09\)70129-6](http://dx.doi.org/10.1016/S0049-237X(09)70129-6).
- [Jac99] Bart Jacobs. *Categorical Logic and Type Theory*. Studies in Logic and the Foundations of Mathematics 141. Amsterdam: North Holland, 1999.
- [KHS19] Ambrus Kaposi, Simon Huber, and Christian Sattler. “Gluing for Type Theory”. In: *4th International Conference on Formal Structures for Computation and Deduction (FSCD 2019)*. Ed. by Herman Geuvers. Vol. 131. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss

- Dagstuhl–Leibniz-Zentrum fuer Informatik, 2019, 25:1–25:19. ISBN: 978-3-95977-107-8. DOI: [10.4230/LIPIcs.FSCD.2019.25](https://doi.org/10.4230/LIPIcs.FSCD.2019.25).
- [KKA19] Ambrus Kaposi, András Kovács, and Thorsten Altenkirch. “Constructing quotient inductive-inductive types”. In: *Proceedings of the ACM on Programming Languages* 3.POPL (Jan. 2019), 2:1–2:24. DOI: [10.1145/3290315](https://doi.org/10.1145/3290315).
- [KL21] Krzysztof Kapulkin and Peter LeFanu Lumsdaine. “The simplicial model of Univalent Foundations (after Voevodsky)”. In: *Journal of the European Mathematical Society* 23 (6 2021), pp. 2071–2126. DOI: [10.4171/JEMS/1050](https://doi.org/10.4171/JEMS/1050).
- [KL23] Chris Kapulkin and Yufeng Li. *Extensional concepts in intensional type theory, revisited*. Preprint. Oct. 2023. arXiv: [2310.05706](https://arxiv.org/abs/2310.05706) [math.LO].
- [Kle50] S. C. Kleene. “A symmetric form of Gödel’s theorem”. In: *Ind. Math* (1950), 12:244–246. DOI: [10.2307/2266709](https://doi.org/10.2307/2266709).
- [Kov] András Kovács. *elaboration-zoo*. Github. URL: <https://github.com/AndrasKovacs/elaboration-zoo/tree/master>.
- [Kov22] András Kovács. “Type-Theoretic Signatures for Algebraic Theories and Inductive Types”. PhD thesis. Eötvös Loránd University, 2022. DOI: [10.15476/ELTE.2022.070](https://doi.org/10.15476/ELTE.2022.070).
- [KS15] Nicolai Kraus and Christian Sattler. “Higher Homotopies in a Hierarchy of Univalent Universes”. In: *ACM Transactions on Computational Logic* 16.2 (Apr. 2015), 18:1–18:12. ISSN: 1529-3785. DOI: [10.1145/2729979](https://doi.org/10.1145/2729979).
- [LH12] Daniel R. Licata and Robert Harper. “Canonicity for 2-Dimensional Type Theory”. In: *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL ’12. New York, NY, USA: ACM, 2012, pp. 337–348. ISBN: 978-1-4503-1083-3. DOI: [10.1145/2103656.2103697](https://doi.org/10.1145/2103656.2103697).
- [Lic16] Dan Licata. *Weak univalence with “beta” implies full univalence*. Email to Homotopy Type Theory mailing list. Sept. 2016. URL: <https://groups.google.com/d/msg/homotopytypetheory/j2KBIVDw53s/YTDK4D0NFQAJ>.
- [LMS10] Andres Löh, Conor McBride, and Wouter Swierstra. “A Tutorial Implementation of a Dependently Typed Lambda Calculus”. In: *Fundam. Informaticae* 102.2 (2010), pp. 177–207. DOI: [10.3233/FI-2010-304](https://doi.org/10.3233/FI-2010-304).
- [LS13] Daniel R. Licata and Michael Shulman. “Calculating the Fundamental Group of the Circle in Homotopy Type Theory”. In: *Proceedings of the 2013 28th Annual ACM/IEEE Symposium on Logic in Computer Science*. LICS ’13. Washington, DC, USA: IEEE Computer Society, 2013, pp. 223–232. ISBN: 978-0-7695-5020-6. DOI: [10.1109/LICS.2013.28](https://doi.org/10.1109/LICS.2013.28).

- [LS88] Joachim Lambek and Philip J. Scott. *Introduction to Higher-Order Categorical Logic*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1988. ISBN: 9780521356534.
- [Lur09] Jacob Lurie. *Higher Topos Theory*. Princeton University Press, 2009. ISBN: 9780691140490.
- [LW15] Peter LeFanu Lumsdaine and Michael A. Warren. “The Local Universes Model: An Overlooked Coherence Construction for Dependent Type Theories”. In: *ACM Transactions on Computational Logic* 16.3 (2015). DOI: 10.1145/2754931.
- [McB02] Conor McBride. “Elimination with a Motive”. In: *Types for Proofs and Programs*. Springer Berlin Heidelberg, 2002, pp. 197–216. ISBN: 9783540458425. DOI: 10.1007/3-540-45842-5_13.
- [McB18] Conor McBride. *Basics of bidirectionality*. Online. 2018. URL: <https://pigworker.wordpress.com/2018/08/06/basics-of-bidirectionality/>.
- [McB19] Conor McBride. *The types who say ‘ni’*. Online. 2019. URL: <https://github.com/pigworker/TypesWhoSayNi/blob/master/tex/TypesWhoSayNi.pdf>.
- [McB99] Conor McBride. “Dependently Typed Functional Programs and their Proofs”. PhD thesis. University of Edinburgh, 1999. URL: <https://era.ed.ac.uk/bitstream/id/600/ECS-LFCS-00-419.pdf>.
- [Mim20] Samuel Mimram. *PROGRAM = PROOF*. Independently published, 2020. ISBN: 979-8615591839. URL: <http://www.lix.polytechnique.fr/Labo/Samuel.Mimram/teaching/INF551/course.pdf>.
- [ML71] Per Martin-Löf. “An intuitionistic theory of types”. Unpublished preprint. 1971.
- [ML75] Per Martin-Löf. “An Intuitionistic Theory of Types: Predicative Part”. In: *Logic Colloquium ’73*. Ed. by H. E. Rose and J. C. Shepherdson. Vol. 80. Studies in Logic and the Foundations of Mathematics. North-Holland, 1975, pp. 73–118. DOI: 10.1016/S0049-237X(08)71945-1.
- [ML82] Per Martin-Löf. “Constructive mathematics and computer programming”. In: *Logic, Methodology and Philosophy of Science VI, Proceedings of the Sixth International Congress of Logic, Methodology and Philosophy of Science, Hannover 1979*. Ed. by L. Jonathan Cohen, Jerzy Łoś, Helmut Pfeiffer, and Klaus-Peter Podewski. Vol. 104. Studies in Logic and the Foundations of Mathematics. North-Holland, 1982, pp. 153–175. DOI: 10.1016/S0049-237X(09)70189-2.

- [ML84a] Per Martin-Löf. “Constructive mathematics and computer programming”. In: *Philosophical Transactions of the Royal Society of London A* 312.1522 (Oct. 1984), pp. 501–518. ISSN: 0080-4614. DOI: [10.1098/rsta.1984.0073](https://doi.org/10.1098/rsta.1984.0073).
- [ML84b] Per Martin-Löf. *Intuitionistic type theory. Notes by Giovanni Sambin of a series of lectures given in Padua, June 1980*. Vol. 1. Studies in Proof Theory. Bibliopolis, 1984. ISBN: 88-7088-105-9.
- [ML87] Per Martin-Löf. “Truth of a Proposition, Evidence of a Judgement, Validity of a Proof”. In: *Synthese* 73.3 (1987), pp. 407–420. DOI: [10.1007/bf00484985](https://doi.org/10.1007/bf00484985).
- [ML92] Per Martin-Löf. *Substitution calculus*. Notes from a lecture given in Göteborg. 1992.
- [ML96] Per Martin-Löf. “On the meanings of the logical constants and the justifications of the logical laws”. In: *Nordic journal of philosophical logic* 1.1 (1996), pp. 11–60.
- [MS93] John C. Mitchell and Andre Scedrov. “Notes on scoping and relators”. In: *Computer Science Logic*. Ed. by E. Börger, G. Jäger, H. Kleine Büning, S. Martini, and M. M. Richter. Springer Berlin Heidelberg, 1993, pp. 352–378. DOI: [10.1007/3-540-56992-8_21](https://doi.org/10.1007/3-540-56992-8_21).
- [MU21] Leonardo de Moura and Sebastian Ullrich. “The Lean 4 Theorem Prover and Programming Language”. In: *Automated Deduction – CADE 28*. Ed. by André Platzer and Geoff Sutcliffe. Cham: Springer International Publishing, 2021, pp. 625–635. ISBN: 978-3-030-79876-5. DOI: [10.1007/978-3-030-79876-5_37](https://doi.org/10.1007/978-3-030-79876-5_37).
- [NFS12] Fredrik Nordvall Forsberg and Anton Setzer. “A finite axiomatisation of inductive-inductive definitions”. In: *Logic, Construction, Computation*. Ed. by Ulrich Berger, Diener Hannes, Peter Schuster, and Monika Seisenberger. Vol. 3. Ontos mathematical logic. Ontos Verlag, 2012, pp. 259–287. DOI: [10.1515/9783110324921.259](https://doi.org/10.1515/9783110324921.259).
- [NPS90] Bengt Nordström, Kent Petersson, and Jan Smith. *Programming in Martin-Löf’s Type Theory*. Oxford University Press, 1990. URL: <http://www.cse.chalmers.se/research/group/logic/book/>.
- [NPW02] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL: A Proof Assistant for Higher-Order Logic*. Vol. 2283. Lecture Notes in Computer Science. Springer, 2002. DOI: [10.1007/3-540-45949-9](https://doi.org/10.1007/3-540-45949-9).
- [Oos08] Jaap van Oosten. *Realizability: An Introduction to its Categorical Side*. Vol. 152. Studies in Logic and the Foundations of Mathematics. Elsevier Science, 2008. ISBN: 9780444515841.

- [PD01] Frank Pfenning and Rowan Davies. “A judgmental reconstruction of modal logic”. In: *Mathematical Structures in Computer Science* 11.4 (2001), pp. 511–540. DOI: [10.1017/S0960129501003322](https://doi.org/10.1017/S0960129501003322).
- [PM93] Christine Paulin-Mohring. “Inductive definitions in the system Coq rules and properties”. In: *Typed Lambda Calculi and Applications*. Springer-Verlag, 1993, pp. 328–345. ISBN: 3540565175. DOI: [10.1007/bfb0037116](https://doi.org/10.1007/bfb0037116). URL: <http://dx.doi.org/10.1007/BFb0037116>.
- [Pol02] Robert Pollack. “Dependently Typed Records in Type Theory”. In: *Formal Aspects of Computing* 13.3–5 (July 2002), pp. 386–402. ISSN: 0934-5043. DOI: [10.1007/s001650200018](https://doi.org/10.1007/s001650200018).
- [PPM90] Frank Pfenning and Christine Paulin-Mohring. “Inductively defined types in the Calculus of Constructions”. In: *Mathematical Foundations of Programming Semantics*. Springer-Verlag, 1990, pp. 209–228. ISBN: 9780387348087. DOI: [10.1007/bfb0040259](https://doi.org/10.1007/bfb0040259).
- [PT00] Benjamin C. Pierce and David N. Turner. “Local type inference”. In: *ACM Transactions on Programming Languages and Systems* 22.1 (2000), pp. 1–44.
- [PT22] Loïc Pujet and Nicolas Tabareau. “Observational Equality: Now for Good”. In: *Proceedings of the ACM on Programming Languages* 6.POPL (Jan. 2022). DOI: [10.1145/3498693](https://doi.org/10.1145/3498693).
- [Ras21] Nima Rasekh. *Univalence in Higher Category Theory*. 2021. arXiv: [2103.12762](https://arxiv.org/abs/2103.12762) [math.CT].
- [Rez10] Charles Rezk. “Toposes and Homotopy Toposes”. Originally circulated in 2005. 2010. URL: <https://rezk.web.illinois.edu/homotopy-topos-sketch.pdf>.
- [Rij+21] Egbert Rijke, Elisabeth Stenholm, Jonathan Prieto-Cubides, Fredrik Bakke, and others. *The agda-unimath library*. 2021. URL: <https://github.com/UniMath/agda-unimath/>.
- [Rij22] Egbert Rijke. *Introduction to Homotopy Type Theory*. To be published. Cambridge University Press, 2022. arXiv: [2212.11082](https://arxiv.org/abs/2212.11082).
- [Ros36] Barkley Rosser. “Extensions of Some Theorems of Gödel and Church”. In: *The Journal of Symbolic Logic* 1.3 (1936), pp. 87–91. ISSN: 00224812. URL: <http://www.jstor.org/stable/2269028> (visited on 08/08/2022).
- [SA21] Jonathan Sterling and Carlo Angiuli. “Normalization for Cubical Type Theory”. In: *36th Annual ACM/IEEE Symposium on Logic in Computer Science*. LICS ’21. 2021, pp. 1–15. DOI: [10.1109/LICS52264.2021.9470719](https://doi.org/10.1109/LICS52264.2021.9470719).

- [SAG22] Jonathan Sterling, Carlo Angiuli, and Daniel Gratzer. “A Cubical Language for Bishop Sets”. In: *Logical Methods in Computer Science* 18 (1 Mar. 2022). DOI: [10.46298/lmcs-18\(1:43\)2022](https://doi.org/10.46298/lmcs-18(1:43)2022).
- [Sco18] Dana Scott. *Looking backwards; Looking forwards*. Invited Talk at the Workshop in honour of Dana Scott’s 85th birthday and 50 years of domain theory. July 2018.
- [Ser53] Jean-Pierre Serre. “Cohomologie modulo 2 des complexes d’Eilenberg-MacLane”. In: *Commentarii Mathematici Helvetici* 27.1 (Dec. 1953), pp. 198–232. ISSN: 1420-8946. DOI: [10.1007/bf02564562](https://doi.org/10.1007/bf02564562).
- [SH06] Christopher A. Stone and Robert Harper. “Extensional equivalence and singleton types”. In: *Transactions on Computational Logic* 7.4 (2006), pp. 676–722. DOI: [10.1145/1183278.1183281](https://doi.org/10.1145/1183278.1183281).
- [Shu08] Michael A. Shulman. *Set theory for category theory*. 2008. arXiv: [0810.1279](https://arxiv.org/abs/0810.1279).
- [Shu19] Michael Shulman. *All $(\infty, 1)$ -toposes have strict univalent universes*. Preprint. Apr. 2019. arXiv: [1904.07004](https://arxiv.org/abs/1904.07004) [math.AT].
- [Shu21] Michael Shulman. “Homotopy Type Theory: The Logic of Space”. In: *New Spaces in Mathematics: Formal and Conceptual Reflections*. Ed. by Mathieu Anel and Gabriel Catren. Vol. 1. Cambridge University Press, 2021. Chap. 6, pp. 322–404. DOI: [10.1017/9781108854429.009](https://doi.org/10.1017/9781108854429.009).
- [Shu23] Michael Shulman. *Towards third generation HoTT*. Joint work with Thorsten Altenkirch and Ambrus Kaposi. Slides available at <https://home.sandiego.edu/~shulman/papers/hott-cmu-day1.pdf>. 2023.
- [Smi88] Jan M. Smith. “The Independence of Peano’s Fourth Axiom from Martin-Löf’s Type Theory Without Universes”. In: *The Journal of Symbolic Logic* 53.3 (1988), pp. 840–845. ISSN: 00224812. DOI: [10.2307/2274575](https://doi.org/10.2307/2274575).
- [SP94] Paula Severi and Erik Poll. “Pure Type Systems with Definitions”. In: *Proceedings of the Third International Symposium on Logical Foundations of Computer Science*. LFCS ’94. Berlin, Heidelberg: Springer-Verlag, 1994, pp. 316–328. ISBN: 3540581405.
- [Ste19] Jonathan Sterling. *Algebraic Type Theory and Universe Hierarchies*. 2019. arXiv: [1902.08848](https://arxiv.org/abs/1902.08848) [cs.LO].
- [Ste21] Jonathan Sterling. “First Steps in Synthetic Tait Computability: The Objective Metatheory of Cubical Type Theory”. PhD thesis. Carnegie Mellon University, 2021. DOI: [10.5281/zenodo.5709838](https://doi.org/10.5281/zenodo.5709838).

- [Str93] Thomas Streicher. “Investigations Into Intensional Type Theory”. Habilitation thesis. Ludwig-Maximilians-Universität München, 1993. URL: <https://www2.mathematik.tu-darmstadt.de/~streicher/HabilStreicher.pdf>.
- [Stu16] Aaron Stump. *Verified Functional Programming in Agda*. Association for Computing Machinery and Morgan & Claypool, 2016. ISBN: 9781970001273. DOI: 10.1145/2841316.
- [Tas93] Álvaro Tasistro. *Formulation of Martin-Löf’s theory of types with explicit substitutions*. Licentiate thesis, Chalmers University of Technology and University of Göteborg, 1993.
- [Tra53] B. A. Trahtenbrot. “On Recursive Separability”. In: *Dokl. Acad. Nauk* 88 (1953), pp. 953–956.
- [Tse17] Dimitris Tsementzis. “Univalent foundations as structuralist foundations”. In: *Synthese* 194.9 (2017), pp. 3583–3617. DOI: 10.1007/s11229-016-1109-x.
- [Tur89] David Turner. “A new formulation of constructive type theory”. In: *Proceedings of the Workshop on Programming Logic*. 1989.
- [Uem21] Taichi Uemura. “Abstract and Concrete Type Theories”. PhD thesis. Institute for Logic, Language and Computation, University of Amsterdam, 2021. URL: <https://hdl.handle.net/11245.1/41ff0b60-64d4-4003-8182-c244a9afab3b>.
- [UF13] The Univalent Foundations Program. *Homotopy Type Theory: Univalent Foundations of Mathematics*. Self-published, 2013. URL: <https://homotopytypetheory.org/book/>.
- [VAG+20] Vladimir Voevodsky, Benedikt Ahrens, Daniel Grayson, et al. *UniMath — a computer-checked library of univalent mathematics*. 2020. URL: <https://github.com/UniMath/UniMath>.
- [Vaz+14] Niki Vazou, Eric L. Seidel, Ranjit Jhala, Dimitrios Vytiniotis, and Simon Peyton-Jones. “Refinement Types for Haskell”. In: *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming*. ICFP ’14. New York, NY, USA: ACM, 2014, pp. 269–282. ISBN: 9781450328739. DOI: 10.1145/2628136.2628161.
- [VD18] Floris Van Doorn. “On the Formalization of Higher Inductive Types and Synthetic Homotopy Theory”. PhD thesis. Carnegie Mellon University, 2018.
- [Voe06] Vladimir Voevodsky. “A very short note on the homotopy λ -calculus”. Unpublished. 2006.

- [Voe10] Vladimir Voevodsky. *The equivalence axiom and univalent models of type theory*. Notes from a talk at Carnegie Mellon University. Feb. 2010. URL: http://www.math.ias.edu/vladimir/files/CMU_talk.pdf.
- [War08] Michael A. Warren. “Homotopy theoretic aspects of constructive type theory”. PhD thesis. Carnegie Mellon University, Aug. 2008. URL: <http://mawarren.net/papers/phd.pdf>.
- [Wer97] Benjamin Werner. “Sets in types, types in sets”. In: *Theoretical Aspects of Computer Software*. Springer Berlin Heidelberg, 1997, pp. 530–546. ISBN: 9783540695301. DOI: [10.1007/bfb0014566](https://doi.org/10.1007/bfb0014566).
- [Win20] Théo Winterhalter. “Formalisation and Meta-Theory of Type Theory”. PhD thesis. Université de Nantes, 2020. URL: <https://theowinterhalter.github.io/#phd>.
- [WKS22] Philip Wadler, Wen Kokke, and Jeremy G. Siek. *Programming Language Foundations in Agda*. Aug. 2022. URL: <https://plfa.inf.ed.ac.uk/22.08/>.
- [Xi07] Hongwei Xi. “Dependent ML: An approach to practical programming with dependent types”. In: *Journal of Functional Programming* 17.2 (2007), pp. 215–286. DOI: [10.1017/S0956796806006216](https://doi.org/10.1017/S0956796806006216).