

DESIGN AND IMPLEMENTATION OF FAMILY POLYMORPHISM FOR INTERACTIVE THEOREM PROVING

by

Ende Jin

A thesis
presented to the University Of Waterloo
in fulfilment of the
thesis requirement for the degree of
Master of Mathematics
in
Computer Science

Waterloo, Ontario, Canada, 2023

© Ende Jin 2023

AUTHOR'S DECLARATION

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

STATEMENT OF CONTRIBUTIONS

Ende Jin was the sole author for Section 6, which were written under the supervision of Dr. Yizhou Zhang.

Jin E, Amin N, Zhang Y. Extensible Metatheory Mechanization via Family Polymorphism: Technical Report. Technical Report CS-2023-01. School of Computer Science, University of Waterloo. <https://cs.uwaterloo.ca/sites/ca.computer-science/files/uploads/files/cs-2023-01.pdf>; 2023.

This thesis consists in part of a manuscript written for publication. The exception to sole authorship of material is as follows:

Research presented in the rest of the thesis: The research was conducted at the University of Waterloo by Ende Jin under the supervision of Dr. Yizhou Zhang. Ende Jin proposed the primary design of the language, prototyped the Coq plugin for realizing the language, implemented and evaluated the case studies, and formalized and proved meta-theoretic arguments for the core type theory. Ende Jin established the complete language design with consultations from Dr. Yizhou Zhang.

Jin E, Amin N, Zhang Y. Extensible Metatheory Mechanization via Family Polymorphism. Proc. of the ACM on Programming Languages (PACMPL). 2023;7. DOI: <https://doi.org/10.1145/3591286>

ABSTRACT

With the growing practice of mechanizing language metatheories, it has become ever more pressing that interactive theorem provers make it easy to write reusable, extensible code and proofs. This thesis presents a novel language design geared towards extensible metatheory mechanization in a proof assistant. The new design achieves reuse and extensibility via a form of family polymorphism, an object-oriented idea, that allows code and proofs to be polymorphic to their enclosing families. Our development addresses technical challenges that arise from the underlying language of a proof assistant being simultaneously functional, dependently typed, a logic, and an interactive tool. Our results include (1) a prototypical implementation of the language design as a Coq plugin, (2) a dependent type theory capturing the essence of the language mechanism and its consistency and canonicity results, and (3) case studies showing how the new expressiveness naturally addresses real programming challenges in metatheory mechanization.

ACKNOWLEDGMENTS

I am grateful for many useful comments and suggestions. I thank Benjamin Delaware, Anastasiya Kravchuk-Kirilyuk, Yao Li, Andrew Myers, Zhixuan Yang, Prabhakar Ragde, Werner Dietl for valuable discussions and feedback.

I also thank my advisors: Yizhou Zhang and Ondřej Lhoták for their patience and guidance.

This work was supported in part by an Amazon research award on automated reasoning. The views and opinions expressed are those of the author and do not necessarily reflect the position of any funding agency.

TABLE OF CONTENTS

Author’s Declaration	ii
Statement of contributions	iii
Abstract	iv
Acknowledgments	v
List of Figures	vii
List of Abbreviations	viii
1 Introduction	1
2 Design Requirements and Challenges	3
3 Language Design	5
3.1 Extensible Inductive Types and Exhaustive Recursion/Induction	5
3.2 Late Binding and Equalities	7
3.3 Overriding	8
3.4 Sound Logical Reasoning	8
3.5 Composing Families as Mixins	8
3.6 Injectivity and Disjointness of Constructors via Partial Recursors	9
4 Compiling Family Polymorphism to Parameterized Modules	14
4.1 Limitations	17
5 A Preview of FMLTT: A Core Dependent Type Theory	18
6 Syntax and Semantic Models of FMLTT	25
6.1 MLTT with Explicit Substitutions and Universe Levels	25
6.2 FMLTT	28
6.3 A Translation that Compiles Linkages Away	31
6.4 A Proof Relevant Logical-Relations Model for Canonicity	32
6.5 Using FMLTT’s Linkage Transformers to Model a Derived Family	38
7 Case Studies	39
8 Related Work	41
9 Conclusion	43
References	44

LIST OF FIGURES

1	STLC metatheories (left) and its extension with fixpoints (right).	3
2	Using FPOP to mechanize STLC and the fixpoints extension, as envisioned in Figure 1.	6
3	Composing extensions of STLC.	9
4	Compilation of family STLC (Figure 2).	15
5	Compilation of family STLC _{fix} and the final Check command (Figure 2).	16
6	Syntax and selected typing rules of MLTT, named variables and meta-level substitution	18
7	Syntax and selected typing rules of FMLTT, named variables and meta-level substitution	19
8	FMLTT encoding of the STLC family from Figure 2	21

LIST OF ABBREVIATIONS

EP Expression Problem	1
OO object-oriented	1
STLC simply typed λ -calculus	1
FPOP <u>F</u> amily <u>P</u> olynomialism for a <u>P</u> roof assistant	5
FMLTT <u>FaMiLy</u> <u>T</u> ype <u>T</u> heory	18
QIIT Quotient Inductive-Inductive Type	23
DTC data types à la carte	41
MTC meta-theory à la carte	41

1 INTRODUCTION

There is a growing trend among programming languages researchers to use proof assistants to mechanize metatheories. However, the programmer runs into an old problem in the new setting of proof engineering: the Expression Problem (EP) [Wadler et al. 1998].

The EP is a programming challenge that epitomizes the difficulty of writing type-safe, extensible code. To define an expression language that can be reused for future extensions, the programmer faces a fundamental tension [Reynolds 1975] between adding new constructors to a data type (e.g., new abstract syntax) and adding new functions over the data type (e.g., new compiler passes).

The EP is well studied in the setting of ordinary functional programming and object-oriented (OO) programming. Modern languages, such as Scala [Odersky and Zenger 2005], have a good supply of linguistic features that offer expressive power to solve the EP.

In contrast, proof assistants offer few linguistic solutions to the EP. But the challenge of writing extensible, type-safe code is as real as in any other language. Metatheory mechanization epitomizes the difficulty: the programmer faces a tension between adding new constructors to an inductive type (e.g., new abstract syntax) and adding new functions and theorems over the inductive type (e.g., new metatheoretical results).

In the Coq proof assistant, for instance, inductive types, as well as functions and theorems over inductive types, are closed to extension. So to reuse mechanized metatheories, the common practice is still to copy code and proofs and then modify them in each extension. But having to maintain multiple copies is highly non-modular and antithetical to good software engineering. The programmer could also turn to design patterns for structuring developments and tool support for cutting down on boilerplate [Delaware et al. 2011, 2013a; Schwaab and Siek 2013; Keuchel and Schrijvers 2013; Forster and Stark 2020]. But these solutions tend to require heavy lifting from the programmer to make code extensible and often lead to non-idiomatic programming styles.


We thus set out to answer the following question: can extensible metatheory mechanization be made easier by having a proof assistant support new *linguistic features* that address the EP?

At the core of many linguistic solutions to the EP is *inheritance*. Inheritance is sometimes interpreted narrowly as a subclass’ inheriting methods and instance variables from its superclass. But the language-theoretic essence of inheritance is more general: it is a linguistic approach to incrementally modifying record-like definitions by allowing *late binding* [Cook et al. 1990].

Language mechanisms including virtual classes [Madsen and Moller-Pedersen 1989; Ernst et al. 2006], mixins [Bracha and Cook 1990], virtual types [Thorup 1997], and extensible cases [Blume et al. 2006] are based on this essential idea of inheritance. In particular, when a language mechanism allows late binding of the meaning of nested types and terms, it is said to support *family polymorphism* [Ernst 2001]: types and terms are polymorphic to a family they are nested within.

Contributions. We contribute a language design that integrates family polymorphism into a proof assistant. Because code and proofs are polymorphic to a family they are nested within, they can be inherited and reused by a derived family. Hence, family polymorphism allows for extensible metatheory mechanization.

As an example, the diagram to the right depicts an extensible mechanization of the simply typed λ -calculus (STLC), using family polymorphism. An extension of STLC with products and another with fixpoints can both inherit from the base STLC family: they reuse mechanized metatheories, from abstract syntax all the way to the type-safety theorem, only adding new constructors to inductive types and adding new cases to recursive functions and induction proofs *as needed* by an extension.



Integrating family polymorphism into a dependent type theory for logical reasoning, however, poses significant technical challenges. As our analysis will show, pillars of dependent type theories—including inductive types, definitional equality, and logical consistency—are all inimical to the kind of extensibility and family polymorphism found in existing OO language designs. Thus, our contributions include novel design recipes for dealing with these challenges and foundational meta-theoretical guarantees on the underlying logic. Specifically, we make the following contributions.

- We present a language design that enables extensible metatheory mechanization in a higher-order, dependent type theory with inductive types (Section 3). The language design reconciles the expressiveness enabled by family polymorphism with the rigor of a proof assistant, while largely retaining an idiomatic programming style.
- We contribute a prototypical implementation of our language mechanism as a Coq plugin (Section 4), available at <https://github.com/DKXXXL/FPOP>. The plugin works by compiling surface-language definitions into Coq modules parameterized by explicit extensibility hooks.
- We capture the essence of the new language mechanism formally by extending Martin-Löf type theory with facilities to express family polymorphism (Section 5). We prove foundational meta-theoretical results including consistency and canonicity.
- We present case studies of using our Coq plugin to mechanize language metatheories (Section 7). They show how our language design naturally solves the EP and enables a good amount of reuse and extensibility for mechanizing proofs.

2 DESIGN REQUIREMENTS AND CHALLENGES

Integrating family polymorphism into a proof assistant presents challenges far beyond those found in an object-oriented setting, as the underlying programming language is simultaneously functional, dependently typed, a logic, and an interactive tool.

C1. Extensible inductive types vs. exhaustive inductive reasoning. Inductive types, generalizing algebraic data types found in functional languages, are a central feature of any proof assistant in use for mechanizing language metatheories. They offer a means to define abstract syntax and inference rules. But unfortunately, inductive types are closed to extension by design.

A family-polymorphism design could potentially support extensible inductive types, by allowing a *derived family* to add new constructors to inductive types inherited from a *base family*. Such a feature would be useful for extending mechanized languages. As an example, Figure 1 shows an STLC extension, the mechanization of which would be made easier by extensible inductive types. Code would be organized into two families, with the derived family inheriting constructors from the base family (e.g., the ellipsis under “Typing rules”) and adding new constructors to model the syntax and semantics of a fixpoint construct.

However, there is a tension between extensibility of inductive types and exhaustivity of inductive reasoning.¹ In Figure 1, all the lemmas and theorems, as well as the substitution function, require induction (i.e., elimination of inductive types). A language design must enforce that induction remains exhaustive in the face of the new constructors in the derived family. For modularity, the

¹The tension reflects a duality between variants and records. With record-like language constructs (e.g., objects and families), it is safe for an extension to add new fields: existing fields can still be projected. But variant-like constructs (e.g., inductive types) do not automatically enjoy safe, modular addition of constructors: existing pattern matches could become non-exhaustive.

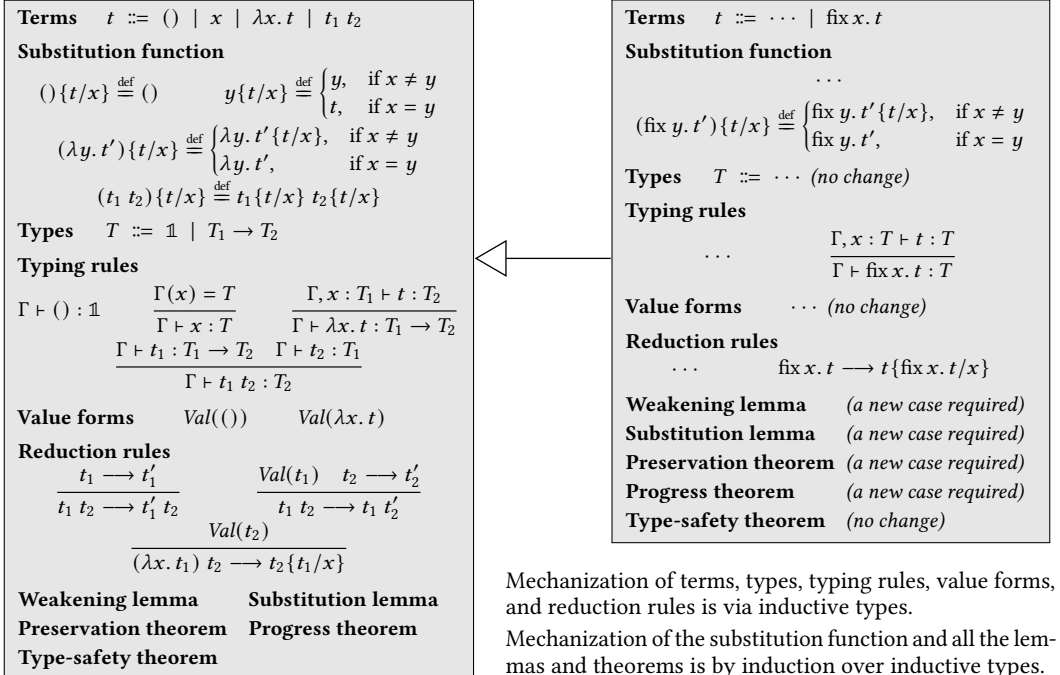


Figure 1. STLC metatheories (left) and its extension with fixpoints (right).

type system should do so without requiring redefinition or rechecking of those cases already handled by the base family.

C2. Late binding vs. definitional equality. Family polymorphism enables modular reuse via late binding: the code of a base family can be reused by a derived family, because fields referenced by that code have meanings polymorphic to the enclosing family.

This flexibility, however, prevents *definitional equality* that one takes for granted when programming in a proof assistant. In Figure 1, a proof assistant supporting family polymorphism cannot unfold references to the substitution function into a pattern match against four cases, as a derived family may modify the definition of the function by adding new cases. Without the ability to unfold the substitution function, how can the programmer even prove the substitution lemma? The problem is compounded by the occasional need in derived families to override fields, which is potentially at odds with being able to use equalities over the fields.

C3. Self reference vs. logical consistency. The language-theoretic essence of late binding is self reference; inheritance and family polymorphism are mechanisms for incrementally modifying self-referential definitions [Cook et al. 1990]. However, self reference could easily lead to divergence. Divergence is not a concern for the design of ordinary OO or functional languages, but it would mean logical inconsistency—and hence unsoundness!—for a language aimed at logical reasoning. A family-polymorphism design must tame self reference to guarantee consistency.

C4. User experience and system implementation. Interactive theorem proving and tactic programming are central to a typical programming experience with a proof assistant. A language design integrating family polymorphism should be compatible with these forms of programming. In particular, it should be possible in our system to incrementally navigate through vernacular commands and, moreover, construct proofs with common tactics while getting instant feedback on the proof state, even in the middle of a family definition. Last but not least, in addition to proving theorems, it should be possible for terms defined with families to possess computational content.

3 LANGUAGE DESIGN

We present the key ingredients of our design as an extension to the Coq proof assistant, though we believe the design could be adapted to other proof assistants such as Lean. We call our design and implementation Family POLymorphism for a Proof assistant (**FPOP**). In this section, we focus on the language design of FPOP. Section 4 describes its implementation as a Coq plugin.

Figure 2 shows how STLC and its extension with fixpoints can be mechanized using FPOP, in a style envisioned in Figure 1. The base family STLC hosts the STLC metatheories, from abstract syntax to the type-safety theorem. Family STLCFix, derived from STLC, makes adjustments as needed by a fixpoints extension: it adds new constructors to the inductive types (**FInductive**) and adds new cases to the recursive functions (**FRecursion**) and induction proofs (**FInduction**). Existing constructors and cases, as well as those definitions and theorems that need no adjustments (ty, env, empty, steps, and typesafe), are automatically inherited and reused. In particular, executing the last command, **Check** STLCFix.typesafe, displays the type-safety theorem of the fixpoints extension.

3.1 Extensible Inductive Types and Exhaustive Recursion/Induction

Extending inductive types. In family STLCFix, the **FInductive** `tm` further binds the `tm` type in family STLC. It has five constructors: four inherited from STLC and a fifth called `tm_fix`.

Crucially, the meanings of `tm` and its constructors are *late bound*, depending on the family in which they are referenced. Consider `tm_app`. It is defined in family STLC and is thus unaware of `tm_fix`. Yet in family STLCFix, we can use `tm_app` to construct applications of fixpoints, as in `tm_app (tm_fix "f" t1) t2`. This use is justified by `tm_app`'s type, $tm \rightarrow tm \rightarrow tm$. It allows `tm_app` to be applied to anything of type `tm`, which in family STLCFix include those constructed from `tm_fix`.

Ensuring exhaustivity of induction. Ordinarily, an inductive type is not extensible: it is exhaustively generated by its constructors and has no more inhabitants beyond those they construct. This idea is captured by the eliminator (aka *recursor*) associated with an inductive type. For example, if `tm` were defined as an ordinary inductive type, then its eliminator would have the following type:

$$\begin{aligned} \text{tm_rect} : & \forall (P : tm \rightarrow \text{Type}), P \text{ tm_unit} \rightarrow (\forall x, P (tm_var x)) \rightarrow \\ & (\forall x t, P t \rightarrow P (tm_abs x t)) \rightarrow \\ & (\forall t1, P t1 \rightarrow \forall t2, P t2 \rightarrow P (tm_app t1 t2)) \rightarrow \forall t, P t \end{aligned}$$

The eliminator would enable function definitions by recursion and proofs by induction, over `tm`, that exhaustively handle the four cases corresponding to each constructor. The (dependent) return type `P` is called the *motive* of the recursion.

With inductive types made extensible, exhaustivity of induction is now in question, however. In particular, the recursor `tm_rect` should no longer be allowed, because its type mentions `tm`, whose meaning is late bound, yet `tm_rect` purports to claim $\forall t, P t$ given only four case handlers.

To reconcile the tension without requiring redefinition or rechecking of case handlers (**C1**), our design introduces the **FRecursion** and **FInduction** commands. The key idea is to allow case handlers to be added retroactively, should inductive types be extended, and to allow recursion and induction (which are defined in terms of case handlers) to be late bound.

As an example, consider the substitution function `subst`, defined using **FRecursion**. The `on` clause specifies that recursion is over `tm`. The `motive` clause suggests that the recursive function being defined has type $tm \rightarrow id \rightarrow tm \rightarrow tm$. The `subst` function in STLC is further bound by the `subst` in STLCFix: the four cases from STLC are automatically inherited and reused, with STLCFix adding a fifth case retroactively to form a new `subst` function. For exhaustivity, it is a static error if the programmer fails to further bind `subst` and define this fifth case. The type system does this check in family STLCFix by examining if the inductive type `tm`, over which `subst` is recursively defined, is further bound in the same family.

```

Family STLC.                                     (* The base STLC *)

FInductive tm : Set :=                           (* Terms *)
| tm_unit : tm
| tm_var   : id → tm
| tm_abs   : id → tm → tm
| tm_app   : tm → tm → tm.

FRecursion subst on tm (* Substitution function *)
  motive λ(_ : tm), id → tm → tm.
Case tm_unit := λ x t, tm_unit.
Case tm_var := λ y x t,
  if (eqb x y) then t else (tm_var y).
Case tm_abs := λ y t' IHT' x t,
  tm_abs y (if (eqb x y) then t' else IHT' x t).
Case tm_app := λ t1 IHT1 t2 IHT2 x t,
  tm_app (IHT1 x t) (IHT2 x t).
End subst.

FInductive ty : Set :=                           (* Types *)
| ty_unit : ty
| ty_arrow : ty → ty → ty.

FDefinition env : Type := id → option ty.
FDefinition empty : env := λ _, None.

FInductive hasty : env → tm → ty → Prop :=
| ht_unit : ∀ G, hasty G tm_unit ty_unit
| ht_var   : ...                               (* Typing rules *)
| ht_abs   : ...
| ht_app   : ...

FInductive value : tm → Prop := (* Value forms *)
| v_unit : value tm_unit
| v_abs  : ∀ x t, value (tm_abs x t).

FInductive step : tm → tm → Prop :=
| st_app1 : ...                               (* Reduction rules *)
| st_app2 : ...
| st_beta : ∀ x t v, value v →
  step (tm_app (tm_abs x t) v) (subst t x v).

FDefinition steps := clos_refl_trans step.

FInduction weakenlem on hasty (* Weaken. lemma *)
  motive λ G t T (_ : hasty G t T),
  ∀ G', includedin G G' → hasty G' t T.
Case ht_unit. ... Qed.      Case ht_var. ... Qed.
Case ht_abs. ... Qed.      Case ht_app. ... Qed.
End weakenlem.

FInduction substlem on hasty (* Subst. lemma *)
  motive λ G' t T (_ : hasty G' t T),
  ∀ G x t' T', G' = extend G x T' →
  hasty empty t' T' → hasty G (subst t x t') T.
Case ht_unit. ... Qed.      Case ht_var. ... Qed.
Case ht_abs. ... Qed.      Case ht_app. ... Qed.
End substlem.

FInduction preserve on hasty (* Preserv. theorem *)
  motive λ G t T (_ : hasty G t T),
  G = empty → ∀ t', step t t' →
  hasty empty t' T.
Case ht_unit. ... Qed.      Case ht_var. ... Qed.
Case ht_abs. ... Qed.      Case ht_app. ... Qed.
End preserve.

FInduction progress on hasty (* Progress theorem *)
  motive λ G t T (_ : hasty G t T),
  G = empty → value t ∨ ∃ t', step t t'.
Case ht_unit. ... Qed.      Case ht_var. ... Qed.
Case ht_abs. ... Qed.      Case ht_app. ... Qed.
End progress.

FTheorem typesafe : (* Type-safety theorem *)
  ∀ t t' T, steps t t' → hasty empty t T →
  value t' ∨ ∃ t'', step t' t''.
Proof. ... Qed.

End STLC.

Family STLCFix extends STLC.
(* STLC extended with fixpoints *)

FInductive tm : Set +=
| tm_fix : id → tm → tm.
FRecursion subst on tm
  motive (λ _, id → tm → tm).
Case tm_fix := λ y t' IHT' x t, ...
End subst.

FInductive hasty : env → tm → ty → Prop +=
| ht_fix : ∀ G x t T, hasty (extend G x T) t T
→ hasty G (tm_fix x t) T.

FInductive step : tm → tm → Prop +=
| st_fix : ∀ x t,
  step (tm_fix x t) (subst t x (tm_fix x t)).

FInduction weakenlem on hasty motive ....
Case ht_fix. ... Qed.
End weakenlem.

FInduction substlem on hasty motive ....
Case ht_fix. ... Qed.
End substlem.

FInduction preserve on hasty motive ....
Case ht_fix. ... Qed.
End preserve.

FInduction progress on hasty motive ....
Case ht_fix. ... Qed.
End progress.

End STLCFix.

| Check STLCFix.typesafe.

```

Figure 2. Using FPOP to mechanize STLC and the fixpoints extension, as envisioned in Figure 1.

The `FInduction` command is similar to `FRecursion` but allows cases to be defined in proof mode. Consider `weakenlem` as an example. It is proven by induction on the typing relation `hasty`. Its `motive` clause shows that the lemma reads as $\forall G\ t\ T, \text{hasty } G\ t\ T \rightarrow \forall G', \text{includedin } G\ G' \rightarrow \text{hasty } G'\ t\ T$. Upon entering case `ht_abs`, for instance, Coq enters proof mode with the goal

$$\forall G', \text{includedin } G\ G' \rightarrow \text{hasty } G'\ (\text{tm_abs } x\ t)\ (\text{ty_arrow } T1\ T2)$$

and with the induction hypothesis $\forall G', \text{includedin } (\text{extend } G\ x\ T1)\ G' \rightarrow \text{hasty } G'\ t\ T2$. The programmer can use tactic programming to discharge the goal. Because the lemma is by `FInduction` on `hasty` and because `STLCFix` adds a constructor `ht_fix` to `hasty`, the programmer is required in family `STLCFix` to extend the proof of `weakenlem` to handle this extra case.

3.2 Late Binding and Equalities

Late binding of nested names. OO inheritance allows the late binding of method names. Family polymorphism generalizes the power of OO inheritance by allowing the late binding of all names nested within families, including those referring to types. We have seen that late binding of `tm` allows the `tm` constructors in `STLC` to be reused in `STLCFix` to construct terms containing `tm_fix`.

As another example, consider the type of `st_beta` in family `STLC`. It refers to `subst`, whose meaning is late bound. When `st_beta` is inherited into family `STLCFix`, `st_beta` has a type that now refers to the `subst` function in `STLCFix`, where `subst` is defined by handling all the five cases known to that family. Thus, late binding of `subst` allows the derived family to reuse `st_beta` as the β -reduction rule for applications possibly constructed from `tm_fix`.

Importantly, a name is late bound only within a family that defines or further binds it. Outside such families, the name can be accessed only by explicitly specifying a family that contains it. For example, the last line of Figure 2 accesses `typesafe` with a qualifier `STLCFix`. The command prints

$$\text{STLCFix.typesafe: } \forall t\ t'\ T, \text{STLCFix.steps } t\ t' \rightarrow \text{STLCFix.hasty } \text{STLCFix.empty } t\ T \rightarrow \text{STLCFix.value } t' \vee \exists t'', \text{STLCFix.step } t'\ t''.$$

In this type, all references to nested names are qualified by `STLCFix`, as desired.

Equality on late bound names. Consider proving the `ht_unit` case of the substitution lemma, `substlem`. The goal is seemingly trivial: the programmer is asked to prove

$$\forall G'\ G\ x\ t'\ T', G' = \text{extend } G\ x\ T' \rightarrow \text{hasty empty } t'\ T' \rightarrow \text{hasty } G\ (\text{subst } \text{tm_unit } x\ t')\ \text{ty_unit}.$$

If `subst` were an ordinary Coq function, then the programmer could discharge the goal with `intros; simpl subst; apply ht_unit`. The Ltac term `simpl subst` unfolds `subst` using its definition and simplifies `subst tm_unit x t'` to `tm_unit`.

But with `subst` being late bound, `subst` cannot and should not be unfolded (C2): the definition of `subst`, as a recursive function, varies across families, yet a derived family should be able to reuse the proof of the `ht_unit` case even when it has to modify the definition of `subst`. Without the ability to unfold `subst`, how can the programmer make progress in this proof, then?

A key observation is that although late binding prevents *definitional equality* on `subst`, it does not affect *propositional equality*. That is, $\forall x\ t, \text{subst } \text{tm_unit } x\ t = \text{tm_unit}$, as a proposition (`Prop`) about the computational behavior of `subst`, should still hold. After all, how `subst` is defined on `tm_unit` does not vary from a base family to a derived family; what can vary is `subst` itself as a recursive function combining the case handlers.

Based on this insight, `FPOP` automatically generates a propositional equality for each case handler defined within `FRecursion`, making the equalities and the recursive function available as axioms for use by the rest of the current family. `FPOP` also provides a tactic `fsimpl` that enables, for instance, simplifying `subst tm_unit x t'` to `tm_unit`. `fsimpl` works by rewriting applications of the axiomatized recursive function using the axiomatized equalities about its computational behaviors.

Note that the definitional equality on `subst` is available *outside* those families that contain `subst`. Within those families, the meaning of `subst` is late bound (i.e., polymorphic to the enclosing family), so only propositional equality is available. In contrast, outside those families, `subst` is always referenced by specifying a family that contains it, as in `STLC.subst` and `STLCFix.subst`. As far as the type checker is concerned, `tm_unit` and `STLCFix.subst tm_unit x t'` are the same thing—the type checker equates them definitionally by unfolding `STLCFix.subst` and performing normalization.

3.3 Overriding

In an OO language, a subclass can override methods of a superclass. Similar expressivity is useful for mechanizing proofs, too. For example, in a derived family, rather than adding new cases to an induction proof, the programmer may prefer overriding the proof entirely, as we observe in our case studies (Section 7).

Overriding is potentially incompatible with having equalities on late bound names, however. Coq distinguishes *opaque* definitions from *transparent* ones. `FPOP` supports the overriding of opaque definitions, which include most proofs. It is safe to override opaque definitions, because the type checker will never attempt to unfold them. For transparent definitions, the common case is that the programmer does not want to override them. In Figure 2, `env`, `empty`, `subst`, and `steps` are transparent—that is, they are not defined with `Qed`. `FPOP` treats transparent definitions as non-overridable by default. Thus, the definitional equalities on `env`, `empty`, and `steps`, as well as the propositional equality on `subst`, are available to the type checker for type-checking the families.

`FPOP` does allow the overriding of transparent definitions explicitly marked as `Overridable` by the programmer. Overriding is made safe by requiring that when an overridable field is overridden, code whose type checking involves unfolding that field should be overridden too. We expect this feature to be used occasionally.

3.4 Sound Logical Reasoning

In Figure 2, just because the `typesafe` theorem in `STLC` is inherited and reused by `STLCFix`, it does not follow that in `STLCFix` the programmer can use `typesafe` to prove progress. If the programmer did, then they would be committing the logical fallacy of circular reasoning: the proof of progress would depend on `typesafe`, yet the proof of `typesafe` depends on progress. Such circularity would easily lead to logical inconsistency. Consider the following two families, where `B` extends `A`:

<pre>Family A. FLemma f : False. Proof. Admitted. FLemma g : False. Proof. apply f. Qed. End A.</pre>	<pre>Family B extends A. FLemma f : False. Proof. apply g. Qed. End B.</pre>
---	--

`B` overrides lemma `f` by proving it using `g`. Lemma `g` is in turn inherited from family `A`, where it is proven using a late bound reference to lemma `f`. Circularity between `f` and `g` allows proving `False`!

To ensure the soundness of logical reasoning (C3), the type system requires that in a derived family, the context in which a field is defined be preserved from the base family. In `STLC`, progress is in the context of `typesafe`. Per the requirement, this relationship must be preserved into `STLCFix`, which prevents the proof of progress from depending on `typesafe` in `STLCFix`.

Note that the requirement still allows a derived family to introduce new declarations into the context of a field. For example, the left column of Figure 3 shows an extension of `STLC` with iso-recursive types, where `tsubst` is introduced into the context of `hasty`. In the event that `FPOP` cannot infer where the programmer intends to place a new field, an annotation is required.

3.5 Composing Families as Mixins

Families can be readily reused to construct larger extensions that *mix in* [Bracha and Cook 1990] the functionalities of the individual families. A family like `STLCFix` can be viewed as a family-to-family


```

Family STLCIsorec extends STLC.
  (* STLC extended with iso-recursive types *)
FInductive tm : Set +=
| tm_fold : tm → tm | tm_unfold : tm → tm.
FRecursion subst on tm motive ... .. End subst.
FInductive ty : Set +=
| ty_var : id → ty | ty_rec : id → ty → ty.
FRecursion tysubst on ty (* Type-level substitution *)
  motive λ(_ : ty), id → ty → ty.
Case ty_unit := .... Case ty_arrow := ....
Case ty_var := .... Case ty_rec := ....
End tysubst.
FInductive hasty : env → tm → ty → Prop +=
| ht_fold : ∀ G t α T,
  hasty G t (tysubst T α (ty_rec α T)) →
  hasty G (tm_fold t) (ty_rec α T)
| ht_unfold : ....
...
(* Other adjustments *)
End STLCIsorec.

Family STLCFixIsorec extends STLC
using STLCFix, STLCIsorec. (* STLC extended
  with fixpoints and iso-recursive types *)
End STLCFixIsorec.

Family STLCProd extends STLC.
  (* STLC extended with products *)
FInductive tm : Set +=
| tm_pair : tm → tm → tm
| tm_fst : tm → tm | tm_snd : tm → tm.
FRecursion subst on tm motive ... .. End subst.
FInductive ty : Set +=
| ty_prod : ty → ty → ty.
...
(* Other adjustments *)
End STLCProd.

Family STLCProdIsorec extends STLC
using STLCProd, STLCIsorec. (* STLC extended
  with products and iso-recursive types *)
FRecursion tysubst on ty
  motive λ(_ : ty), id → ty → ty.
Case ty_prod := .... (* Substitution on product types *)
End tysubst.
End STLCProdIsorec.

Family STLCFixProdIsorec extends STLC
using STLCFix, STLCProdIsorec. (* STLC extended
  with fixpoints, products, and iso-recursive types *)
End STLCFixProdIsorec.

```

Figure 3. Composing extensions of STLC.

functor—and hence a mixin, in the sense of [Flatt et al. 1998]—that transforms any family providing the base STLC functionalities (i.e., STLC or a derived family there of) into a new family additionally supporting fixpoints.

In Figure 3, STLCFixIsorec is declared as an STLC extension that mixes in STLCFix and STLCIsorec. The family is declared with minimal verbiage, yet STLCFixIsorec.typesafe is automatically a proof of the type-safety theorem of an STLC equipped with fixpoints and iso-recursive types.

Mixin composition is a form of multiple inheritance, which may cause name conflicts in general. FPOP requires the programmer to resolve conflicts by overriding conflicted overridable fields.

In the presence of extensible inductive types, mixin composition may also create an obligation to retrofit the mixins with new case handlers. In Figure 3, family STLCProdIsorec is composed of two mixins: STLCProd, which extends the inductive type ty with a new constructor ty_prod, and STLCIsorec, which introduces a function tysubst recursively defined on ty. Hence, for exhaustivity, it is required that a composition of STLCProd and STLCIsorec should additionally handle the ty_prod case in tysubst.

3.6 Injectivity and Disjointness of Constructors via Partial Recursors

Tactics support for constructors. Coq provides tactics for proving injectivity and disjointness of constructors (i.e., `injection` and `discriminate`). The proof terms generated by the tactics involve exhaustively matching on the constructors of an inductive type, so they do not work for extensible inductive types (C1). In principle, the programmer could use `FInduction` to prove injectivity and disjointness. But this workaround is unsatisfying: it is tedious, it forces the programmer to revisit the induction proofs every time an inductive type is extended, and above all, why should a property like $\neg(\text{tm_var } "x" = \text{tm_abs } "y" \ t)$ have anything to do with `tm_fix`?

To provide a streamlined programming experience (C4), FPOP offers two tactics, `finjection` and `fdiscriminate`. For example, in a proof state that contains a manifestly false assumption

$H : \text{tm_var } "x" = \text{tm_abs } "y" \ t$, the programmer can use `fdiscriminate` H to obtain `False` and thus discharge the current goal, just as they would with `discriminate` H if `tm` were not extensible.

Partial recursors. We make the observation that injectivity and disjointness of existing constructors ought to hold regardless of future addition of constructors. This insight motivates the design of *partial recursors*, which power the `finjection` and `fdiscriminate` tactics. Partial recursors can be automatically generated for inductive types defined with `FInductive`.

As analyzed earlier, ordinary recursors, such as `tm_rect`, are impossible within a family in which the name of the inductive type is late bound. However, a key observation is that extensible inductive types still admit a weakened elimination principle where the motive is an `option` type. For example, within family `STLC`, the partial recursor for `tm` has the following type, with extra equations describing its behavior on each constructor:

```
tm_prect_STLC : ∀ (P : tm → Type), option (P tm_unit) → (∀ x, option (P (tm_var x))) →
  (∀ x t, option (P t) → option (P (tm_abs x t))) →
  (∀ t1, option (P t1) → ∀ t2, option (P t2) → option (P (tm_app t1 t2))) → ∀ t, option (P t)
(* For arbitrary P r1 r2 r3 r4, and we denote tmPR := tm_prect_STLC P r1 r2 r3 r4 *)
(* We will have the following equations *)
tm_unit_eq_STLC : tmPR tm_unit = r1
tm_var_eq_STLC : tmPR (tm_var x) = r2 x
tm_abs_eq_STLC : tmPR (tm_abs x b) = r3 a b (tmPR b)
tm_app_eq_STLC : tmPR (tm_app a b) = r4 a (tmPR a) b (tmPR b)
```

`tm_prect_STLC` cannot be a trivial function even though its return type has `option`. Because these four equations uniquely describe the `tm_prect_STLC` on the four constructors and make `tm_prect_STLC` behave similarly to the ordinary dependent eliminator of the non-extensible version of `tm`.

Property about Partial Recursors. A partial recursor is simply adding partiality into the signature of the ordinary eliminator—by wrapping each motive with `option`—together with the equations describing its computational behaviors inherited from the ordinary eliminators.

This partial recursor seems non-canonical compared to the (dependent) eliminator for inductive types as the latter can be considered a universal property. This extensionality makes dependent eliminator one of *the most powerful* reasoning tools around inductive type.

Even though we cannot prove partial recursors are *the most powerful* reasoning tool for extensible inductive types, the fact that propositional partial recursors can derive injectivity already demonstrates itself as a good *extensional characterization* of (non-indexed) extensible inductive type (under certain constraints)—because we can “embed” vanilla inductive types into types supporting a partial recursor: a bit more formally speaking, and restricting our focus to non-indexed inductive types:

THEOREM 3.1. *Given a list of n pairs of types $\{x : A_i \vdash B_i(x)\}_i$ s.t.*

- (Detectable Partiality) for each pair $x : A_i \vdash B_i(x)$,
 - Axiom** detectable:
 - forall** $\{T\} \{a\} (f : B_i a \rightarrow \text{option } T)$,
 - $\{\text{forall } x, f x \lt \text{None}\} + \{\text{exists } x, f x = \text{None}\}$.
- Define C as the inductive type using this list, i.e.,
 - Inductive** $C : \text{Set} := \dots \mid c_i : \text{forall } (x : A_i), (B_i x \rightarrow C) \rightarrow C \mid \dots$
- Assume an arbitrary type $D : \text{Set}$ with
 - n **functions** $d_i : \text{forall } (x : A_i), (B_i x \rightarrow D) \rightarrow D$
 - a **partial recursor**

```

prec : forall (R : Set), ...,
  (r_i : (forall (x : A_i), (B_i x → option R) → option R)), ...,
  D → option R
– and n (propositional) computational axioms: for all i,
prec {R} r_1 r_2 .. r_n (d_i a_i b_i)
  = r_i a_i (fun x ⇒ prec {R} r_1 r_2 .. r_n (b_i x))

```

Then there exists an **embedding** from C to D. More concretely, we can have $\text{inj} : C \rightarrow D$ and $\text{linv} : D \rightarrow \text{option } C$ defined using the eliminator of C and prec of D (both acting like identity) such that

```
forall c : C, linv (inj c) = some c
```

PROOF. First, for a fixed $x : A_i$, given $f_i : B_i x \rightarrow T$, we have $\text{lift } f_i : B_i x \rightarrow \text{option } T$, simply wrapping option around the returning result. Second, with (**Detectable Partiality**) we can construct

Definition unlift:

```
forall {a} {T : Set} (f: (forall c: B a, option T))
  (h : forall x, f x <> None),
{g : (forall c: B a, T) | f = lift g}.
```

Theorem detect_lifted:

```
forall {a} {T : Set} (f: (forall c: B a, option T)),
{ g | f = lift g } + {exists x, f x = None}.
```

The idea of unlift is that if f never returns None, then it is always a result of a lift. detect_lifted is just another way of saying (**Detectable Partiality**).

Then we can construct inj and linv as follows, using eliminator and partial recursor.

Definition inj : C → D :=

```
C_rect (fun (_ : C) ⇒ D)
...
(fun (x_i : A_i) (_ : B_i x → C_i) (X_i : B_i x → D) ⇒ d_i x_i X_i)
... .
```

Definition linv : D → option C :=

```
prec C
...
(fun (a_i : A_i) (b_i' : B_i a_i → option C) ⇒
  match (detectable b_i') with
  | left h ⇒
    let (b_i, _) := unlift b_i' h in Some (c_i a_i b_i)
  | right _ ⇒ None
end)
... .
```

Since each “constructor” in C and D are well-corresponded, when invoking C_rect and prec , each case is handled similarly, and they only vary along the $\{A_i \vdash B_i\}_i$ types.

The definition of inj and linv is well-formed but their behavior is still buried in this complicated formulation. The following two equations can clarify them.

(* inj_compute can be proved thanks to
the computational behaviour about C_rect built-in Coq *)

```

Lemma inj_compute : forall a b,
  inj (ci a b) = di a (fun x => inj (b x)).

```

(* linv_compute can be proved thanks to
the n (propositional) computational axioms about prec *)

```

Lemma linv_compute : forall a b,
  linv (di a b) =
  let b' := fun x => linv (b x) in
  match detect_lifted b' with
  | inleft (exist _ x _) => Some (ci a x)
  | inright _ => None
  end.

```

Finally, these two equations are enough to prove that linv is really a left inverse, via induction

```

linv (inj (ci a b))
= linv (di a (fun x => inj (b x)))      (* by inj_compute *)
= Some (ci a (fun x => linv (inj (b x)))) (* by linv_compute *)
= Some (ci a (fun x => b x))             (* by I.H. *)
= Some (ci a b)                         (* by eta rule *)

```

□

(Detectable Partiality) is derivable for all the finitary branching inductive type²—that includes all of the examples shown in this paper.

(Detectable Partiality) is not generally supported when using a function (infinitary branching) in the constructors, for example, the case of using *higher order abstract syntax* to represent functions in tm. However, in those cases, we still have the partial recursor provable and useful—we just cannot see the partial recursor as a strong toolkit anymore.

We emphasize that in [Theorem 3.1](#), each function d_i can actually be considered a constructor because they can be reflected to real constructors of C using the left inverse linv. Thus, [Theorem 3.1](#) implies that (1) every future extension of the inductive type can support this partial recursion (trivially); (2) every type supporting this partial recursor with its computational axioms at least supports these constructors because of the embedding. In other words, *a type (at least) supports these constructors if and only if this type supports the corresponding partial recursor (with the computational axioms)*. Thus, we can argue that the partial recursor gives a sound and complete extensional characterization of all the extensions of a given (non-indexed) inductive type.

We can support a partial dependent eliminator as well. Though the partial dependent eliminator can deduce prec, our prec is enough to construct this left inverse for (non-indexed) inductive type.

```

pdelim : ∀ (P : D → Set), ...,
  (ri : (∀ (x : Ai) (w : Bi x → D), (∀ (b : Bi x), option (P (w b))))
  → option (P (di x w))), ..., ∀ (d : D) → option (P d)

```

To generalize to indexed inductive type, we need to base on indexed W type [[Martin-Löf 1982](#); [Morris and Altenkirch 2009](#); [Hugunin 2017](#)]. We have mechanized the argument for indexed W type and dependent eliminator, but for the case with only one constructor in C and D . This mechanization³ is a direct generalization of the proof for [Theorem 3.1](#).

²By simply enumerating through the finite domain B_i a.

³<https://gist.github.com/DKXXXL/9b4e48350b966fa3739cfb71f1a9a852>

As mentioned earlier, `tm_prect_STLC` is axiomatized along with four equalities describing its computational behaviors, one for each constructor. After unfolding the definition of `tmPR`, the equality for constructor `tm_abs` is as follows:

`tm_abs_eq_STLC`: $\forall x\ t\ P\ H1\ H2\ H3\ H4,$
 $\text{tm_prect_STLC } P\ H1\ H2\ H3\ H4\ (\text{tm_abs } x\ t) = H3\ x\ t\ (\text{tm_prect_STLC } P\ H1\ H2\ H3\ H4\ t)$

Importantly, unlike the standard `tm_rect` recursor, the partial recursor `tm_prect_STLC` is compatible with the late binding of `tm` in its type. When `tm_prect_STLC` is inherited into family `STLCFix`, all the previous four equalities still hold, and a trivial, fifth equality is made available:

`tm_fix_eq_STLC` : $\forall x\ t\ P\ H1\ H2\ H3\ H4,$ `tm_prect_STLC` $P\ H1\ H2\ H3\ H4\ (\text{tm_fix } x\ t) = \text{None}$

Partial recursors appear weaker than ordinary recursors, but there is power in restraint. In particular, they offer a principled, uniform way to derive injectivity and disjointness of constructors, while supporting future extension: they enable injective mappings from a late bound inductive type, like `tm`, to an ordinary inductive type, like `ℕ`, the injectivity and disjointness of whose constructors are readily available. For example, in a proof state with the assumption `H : tm_var "x" = tm_abs "y" t`, running `fdiscriminate H` first applies an injective mapping to both sides of `H`, obtaining

`tm_prect_STLC (λ_, ℕ) (λ_, Some 1) (Some 2) (λ____, Some 3) (λ_____, Some 4) (tm_var "x") =`
`tm_prect_STLC (λ_, ℕ) (λ_, Some 1) (Some 2) (λ____, Some 3) (λ_____, Some 4) (tm_abs "y" t)`

and then rewrites the above equality using the axiomatized computational behaviors of `tm_prect_STLC`, obtaining `Some 1 = Some 3`, from which `False` easily follows. Note that the proof term generated by `fdiscriminate H` in `STLC` is reusable by family `STLCFix`, because the partial recursor and its computational behaviors are compatible with the late binding of `tm`.

Within family `STLCFix`, a second partial recursor (called `tm_prect_STLCFix`) and its computational behaviors are automatically axiomatized, which allows properties of `tm_fix`, such as `tm_fix x t1 = tm_fix x t2 → t1 = t2`, to be proved.

4 COMPILING FAMILY POLYMORPHISM TO PARAMETERIZED MODULES

We implement our language design as a Coq plugin. It works by translating programs in FPOP syntax into programs that can be checked and evaluated by Coq. The translation is compatible with interactive theorem proving (C4), in that a family is translated piece by piece, allowing each field to be defined and checked separately. The translation is modular and efficient, in that code compiled for fields of a base family can be shared with derived families without having to be rechecked.

Explicit self parameterization. The spirit of the translation is to take “family polymorphism” literally: every field is translated into a Coq definition that is polymorphic to (i.e., universally quantified over) a representation of its enclosing family. While this universal quantification has been implicit with the FPOP syntax, it has to be made explicit in the translated Coq code.

Figures 4 and 5 illustrate the translation of the STLC and STLCFix families from Figure 2. Fields of a family are translated into parameterized Coq modules (or parameterized module types).

As an example, consider field `env` in family STLC. It is translated into a top-level module named `STLC°env`. This module has a parameter called `self` representing the enclosing family: fields of the current family in the context of `env` can be referenced through `self`. In particular, `env` is defined as `id → option ty`, where `ty` is a late-bound reference to the `ty` field of the enclosing family. Hence, this reference to `ty` is translated to `self.ty`, which is manifestly polymorphic to the enclosing family. This translation of the `env` field can be shared with a derived family even if it extends `ty` (e.g., `STLCProd`)—no recompilation is needed because `self.ty` is not tied to any concrete definition of `ty`.

The type of `STLC°env`’s `self` parameter is `STLC°env°Ctx`, a module type constructed from `STLC°ty` (i.e., the translation of the field before `env`) and its context `STLC°ty°Ctx`. In turn, `STLC°ty°Ctx` (not shown in Figure 4) is constructed from `STLC°subst`, the translation of the field before `ty`, and its context `STLC°subst°Ctx`. Thus, the `self` parameter can be used to reference those and only those fields in the current field’s typing context, which echoes the discussion in Section 3.4.

Translating extensible inductive types. An `FInductive` definition is translated to a parameterized module type. Consider the inductive type `tm`. In Figure 4, it is translated to a top-level module type `STLC°tm` that declares a `tm` type, four functions standing for the constructors, a partial recursor (`tm_prect_STLC`), and the computational behaviors of the partial recursor (e.g., `tm_abs_eq_STLC`).

Importantly, `STLC°tm` merely declares the existence of these names and their types; it does not specify their definitions. Having these names and their types available through the context parameters (`self`) suffices for the translations of the subsequent fields to be type-checked by Coq. Leaving the definitions unspecified enables STLC and STLCFix to instantiate `tm` differently upon `End STLC` and upon `End STLCFix`. In particular, non-exhaustive pattern matching is prevented because an ordinary recursor like `tm_rect` is not available.

The command `FInductive tm : Set += tm_fix : ...` in family STLCFix is again translated to a module type `STLCFix°tm` (Figure 5). It includes all the names declared by `STLC°tm` via command `Include STLC°tm(self)`, and additionally declares `tm_fix`, a partial recursor, and related equalities.

Translating recursion and induction. An `FRecursion` definition is translated in two parts: first a module containing the definitions of all the case handlers, and then a module type declaring the existence of the recursive function as well as its computational behaviors.

Consider the translation of `subst` in family STLC. First, a module named `STLC°subst°Cases` is generated on the fly. Importantly, programming remains interactive, as the programmer need not wait until the entire `FRecursion` definition is completed to have a `Case` command type-checked.

Upon `End subst`, a module type named `STLC°subst` is generated. As discussed in Section 3.2, `subst` can be further bound, so its definition is not exposed to the fields that come after it. Accordingly, the

```

(* Code emitted upon definition of tm in family STLC *)
Module Type STLC°tm°Ctx.
End STLC°tm°Ctx.

Module Type STLC°tm (self : STLC°tm°Ctx).
  Axiom tm : Set.
  Axiom tm_unit : tm.      Axiom tm_var : id → tm.
  Axiom tm_abs : id → tm → tm.
  Axiom tm_app : tm → tm → tm.
  Axiom tm_prect_STLC : ....
  Axiom tm_unit_eq_STLC : .... Axiom tm_abs_eq_STLC : ....
  Axiom tm_abs_eq_STLC : .... Axiom tm_app_eq_STLC : ....
End STLC°tm.

(* Code emitted for definition of subst in family STLC *)
Module Type STLC°subst°Cases°Ctx.
  Include STLC°tm°Ctx.  Include STLC°tm.
End STLC°subst°Cases°Ctx.

Module STLC°subst°Cases (self : STLC°subst°Cases°Ctx).
  Def subst°tm_unit := (* emitted upon definition of case *)
    λ (x : id) (t : self.tm), self.tm_unit.
  Def subst°tm_var := ....
  Def subst°tm_abs := ....  Def subst°tm_app := ....
End STLC°subst°Cases.

Module Type STLC°subst°Ctx.
  Include STLC°subst°Cases°Ctx.
  Include STLC°subst°Cases.
End STLC°subst°Ctx.

Module Type STLC°subst (self : STLC°subst°Ctx).
  Axiom subst : self.tm → id → self.tm → self.tm.
  Axiom subst_tm_unit_eq :
    ∀ x t, self.subst (self.tm_unit) x t =
      self.subst°tm_unit x t.
  Axiom subst_tm_var_eq : ....
  Axiom ...
End STLC°subst.

(* Code emitted upon definition of ty in family STLC *)
...

(* Code emitted upon definition of env in family STLC *)
Module Type STLC°env°Ctx.
  Include STLC°ty°Ctx.  Include STLC°ty.
End STLC°env°Ctx.

Module STLC°env (self : STLC°env°Ctx).
  Def env : Type := id → option self.ty.
End STLC°env.

(* Code emitted for other fields defined in family STLC *)
...

(* Code emitted upon conclusion of family STLC *)
Module STLC.

  (* Instantiate tm & its constructors *)
  Inductive tm : Set :=
  | tm_unit | tm_var (v : id) ...
  (* Inst. partial recursor & computational behaviors *)
  Def tm_prect_STLC P :=
    tm_rect (λt, option (P t)).
  Fact tm_unit_eq_STLC : .... reflexivity. Qed.
  Fact ...

  Include STLC°subst°Cases.
  (* Instantiate subst & its computational behaviors *)
  Def subst := tm_rect _ subst°tm_unit
    subst°tm_var subst°tm_abs subst°tm_app.
  Fact subst_tm_unit_eq : .... reflexivity. Qed.
  Fact ...

  (* Instantiate ty, its constructors, partial recursor, etc. *) ...
  Include STLC°env. (* Include env *)

  (* Include/Instantiate other fields of family STLC *) ...
End STLC.

```

Figure 4. Compilation of family STLC (Figure 2).

translation `STLC°subst` merely declares the types of `subst` and the equalities about its computational behaviors, leaving `subst` undefined and the equalities unproven. The equalities are stated in terms of the case handlers, whose definitions *are* available through the `self` parameter. So Coq can simplify, for example, the type of `subst_tm_unit_eq` to $\forall x t, \text{self.subst } \text{self.tm_unit } x t = \text{self.tm_unit}$. These equalities about the computational behaviors of `subst` will be included and available for use in the translations of the subsequent fields through their `self` parameters.

Importantly, code generated for the case handlers is shared with derived families. In Figure 5, module `STLCfix°subst°Cases` reuses—without rechecking—all the case handlers in `STLC°subst°Cases` via command `Include STLC°subst°Cases(self)`.

The translation of `FInduction` is similar, except that there is no need to register computational behaviors, as `FInduction` proofs are considered opaque.

Translation of further-bindables vs. non-further-bindables. In family STLC, field `env` and the case handlers for `subst` are not further-bindable by derived families. In contrast, `tm`, `subst`, and the related equalities can be further bound. The distinction is reflected in the translations. The further-bindable fields are translated to module types that export only types of the fields. The non-further-bindable fields are translated to modules that export definitional equalities on the

```

(* Code emitted upon definition of tm in STLCFix *)
Module Type STLCFix°tm°Ctx.
End STLCFix°tm°Ctx.

Module Type STLCFix°tm (self : STLCFix°tm°Ctx).
  Include STLC°tm(self).
  Axiom tm_fix : id → tm → tm.
  Axiom tm_fix_eq_STLC : ∀ ..., tm_prect_STLC ... = None.
  Axiom tm_prect_STLCFix : ....
  Axiom tm_fix_eq_STLCFix : ....
End STLCFix°tm.

(* Code emitted upon definition of subst in STLCFix *)
Module Type STLCFix°subst°Cases°Ctx.
  Include STLCFix°tm°Ctx. Include STLCFix°tm.
End STLCFix°subst°Cases°Ctx.

Module STLCFix°subst°Cases
(self : STLCFix°subst°Cases°Ctx).
  Include STLC°subst°Cases(self). (* reuse *)
  Def subst°tm_fix := .... (* translation of new case *)
End STLCFix°subst°Cases.

Module Type STLCFix°subst°Ctx.
  Include STLCFix°subst°Cases°Ctx.
  Include STLCFix°subst°Cases.
End STLCFix°subst°Ctx.

Module Type STLCFix°subst (self : STLCFix°subst°Ctx).
  Include STLC°subst(self).
  Axiom subst_tm_fix_eq : ....
End STLCFix°subst.

(* Code emitted for other fields defined in STLCFix *) ...
(* Code emitted upon conclusion of STLCFix *)
Module STLCFix.

  (* Instantiate tm & its constructors *)
  Inductive tm : Set :=
  | tm_unit | tm_var (v : id)
  | tm_abs (v : id) (b : tm) | tm_app (a b : tm)
  | tm_fix (v : id) (b : tm).
  (* Inst. tm partial recursors & their comp. behaviors *) ...

  Include STLCFix°subst°Cases.
  (* Instantiate subst & its computational behaviors *)
  Def subst := tm_rect _ subst°tm_unit
  subst°tm_var subst°tm_abs subst°tm_app (* reuse *)
  subst°tm_fix.
  Fact subst_tm_unit_eq : .... reflexivity. Qed.
  Fact ...

  (* Include ty, its constructors, partial recursor, etc. *) ...

  Include STLC°env. (* reuse *)

  (* Include/Instantiate other fields of STLCFix *) ...

  Include STLC°typesafe. (* reuse *)
End STLCFix.

(* Code emitted upon command Check STLCFix.typesafe *)
Check STLCFix.typesafe.

```

Figure 5. Compilation of family STLCFix and the final `Check` command (Figure 2).

fields. Opaque fields in FPOP can be further bound (Section 3.3); they are translated to Coq modules that export opaque fields.

Eliminating `self` by aggregation. Upon the conclusion of a family definition, a representation of the family is created. For example, module STLC in Figure 4 is generated upon `End STLC`. This module can be viewed as the “fixed point” of the `self`-parameterized translations. The “fixed point” is taken step by step, by adding the translation of each field to this module in the same order as they appear in the family definition.

For the non-further-bindables, the translated modules are directly included (e.g., `Include STLC°env` and `Include STLC°subst°Cases` in Figure 4). The instantiation of `self`s for these modules is implicit, thanks to a Coq nicety: when including a higher-order module, Coq automatically instantiates its parameter with the current interactive module environment. For instance, command `Include STLC°subst°Cases` is successfully executed, because Coq automatically instantiates the `self` parameter using the current module environment, which by construction contains all the fields required by `STLC°subst°Cases°Ctx`.

For the further-bindables, `Axioms` declared in the module types must be instantiated.

- In Figure 4, an inductive type `tm` is generated, instantiating the axiomatized `tm` type and its constructors. The partial recursor `tm_prect_STLC` is defined with the help of `tm_rect`, the recursor Coq generates for `tm`. The computational behaviors of `tm_prect_STLC` are immediate, by `reflexivity`.
- Similarly, `subst` is instantiated by applying the recursor `tm_rect` to the (already included) case handlers. The computational behaviors of `subst` are then immediate, by `reflexivity`.

Module STLCFix in Figure 5 is emitted upon `End STLCFix`, in the same way as described above for STLC. The translation makes sharing evident. In particular, case handlers compiled for STLC are

reused to instantiate `subst`, `substlem`, and `alike`. `STLCoenv` and `STLCotypesafe` are also reused in the construction of module `STLCFix`. One may argue that since the first four constructors of `tm` are repeated in `STLCFix`, the translation does not satisfy the modular compilation requirement. We could address this concern by using wrapper types, but we consider restating constructors a small price to pay in return for the clarity and concision of implementation. We emphasize that compiled case handlers, such as `substotm_abs`, are entirely reusable without rechecking, even with restated constructors. Finally, the reference `STLCFix.typesafe` (where `STLCFix` is a family) can simply be translated to `STLCFix.typesafe` (where `STLCFix` is a Coq module), as the last line of Figure 5 shows.

Trusted base. Rather than modifying the Coq kernel to extend its core theory, a translation to Coq conveniently reduces the trusted base of any development using `FPOP` to Coq. In particular, once a family is closed, `Print Assumptions` can be used to verify that there are no lingering `Axioms` introduced by the translation. The ramifications of possible bugs in the `FPOP` implementation are limited to the usability of the plugin.

4.1 Limitations

The `FPOP` implementation currently does not yet bring extensibility to Coq's full facility for inductive types. Mutually inductive types and parameterized inductive types are not yet extensible, though indexed inductive types are supported and can be used to encode parameterized ones. Partial recursors are automatically generated only for non-indexed inductive types. These features are not exercised by our case studies (Section 7) but may be useful for modeling other languages. We believe that they do not pose conceptual challenges and can be addressed with more engineering effort on the same level as the current `FPOP` implementation.

What seems to require more thought from a language-design perspective is the restriction that recursion and induction (on extensible types) cannot be nested. A possible solution is to make the plugin generate proof obligations for nested pattern matches when the inductive types acquire new constructors. Also interesting for future research is the support for automatically converting terms to propositionally equally typed forms using generated propositional equalities. The lack of this automation currently may cause inconveniences for developments using intrinsically typed syntax. Finally, work on nested inheritance [Nystrom et al. 2004; Zhang and Myers 2017] points to a direction to further increase the expressive power of our language design.

5 A PREVIEW OF FMLTT: A CORE DEPENDENT TYPE THEORY

We contribute **FaMiLy Type Theory (FMLTT)**, a core type theory that extends **Martin-Löf** dependent type theory (MLTT) with facilities to express family polymorphism while maintaining consistency and canonicity.

For accessibility of the main text, Section 5 presents FMLTT using named variables and meta-level substitution. For clarity of proof details, Section 6 supplementing Section 5 instead uses **de Bruijn** indices and explicit substitution [Abadi et al. 1989]. We acknowledge that the presentation is dense for an audience without intimate knowledge of MLTT, so we summarize the salient points first.

Summary. FMLTT is intended as a foundational model. So unlike our programmer-facing plugin, fields automatically axiomatized by the plugin require explicit definitions in FMLTT. FMLTT provides MLTT-style constructs that can be used to express families and family polymorphism. Most notably, it extends MLTT with what we call *linkages*. Linkages are a namesake of the theoretical device through which **Zhang and Myers [2017]** model family polymorphism in an OO setting, but the technical details differ significantly from the prior work.

- Linkages model families, so they are like tuples composed of fields (with field names represented by variable bindings). But there is a twist: linkages support late binding. Unlike dependent tuples where a later component is *existentially* quantified over the earlier ones, a linkage component is *universally* quantified over—and thus polymorphic to—the components preceding it.
- FMLTT features *linkage transformers*, which model how a family can be inductively constructed by inheriting fields from another family, adding new fields, and overriding existing fields.
- Inductive types are modeled as W-types [Martin-Löf 1984] and their extension as overriding.
- Consistency and canonicity of FMLTT are proved by giving semantic interpretations to the syntactic typing judgments.

Contexts	$\Gamma, \Delta ::= \cdot \mid \Gamma, A$			
Types	$A, B, T ::= \cup \mid \mathbb{B} \mid \perp \mid \top \mid \Pi(x : A).B \mid \Sigma(x : A).B \mid \text{Eq}(t_1, t_2) \mid \mathbb{S}(t) \mid \text{El}(t)$			
Terms	$t ::= x \mid c(T) \mid () \mid \mathbf{tt} \mid \mathbf{ff} \mid \text{if}(t_1, t_2, t_3) \mid \lambda x.t \mid \text{app}(t_1, t_2) \mid (t_1, t_2) \mid \text{fst } t \mid \text{snd } t \mid \text{refl}(t) \mid \mathbf{J}(t_1, t_2)$			
$\boxed{\Gamma \vdash}$	$\boxed{\Gamma \vdash T}$	$\boxed{\Gamma \vdash T_1 \equiv T_2}$	$\boxed{\Gamma \vdash t : T}$	$\boxed{\Gamma \vdash t_1 \equiv t_2 : T}$
(TM/LAM)	(TM/APP)		(TM/PAIR)	
$\frac{\Gamma \vdash A \quad \Gamma, x : A \vdash t : B}{\Gamma \vdash \lambda x.t : \Pi(x : A).B}$	$\frac{\Gamma \vdash t : \Pi(x : A).B \quad \Gamma \vdash t' : A}{\Gamma \vdash \text{app}(t, t') : B[t'/x]}$		$\frac{\Gamma, x : A \vdash B \quad \Gamma \vdash t_1 : A \quad \Gamma \vdash t_2 : B[t_1/x]}{\Gamma \vdash (t_1, t_2) : \Sigma(x : A).B}$	
(TM/FST)	(TM/SND)	(TY/EL)	(TM/C)	(TMEQ/C)
$\frac{\Gamma \vdash t : \Sigma(x : A).B}{\Gamma \vdash \text{fst } t : A}$	$\frac{\Gamma \vdash t : \Sigma(x : A).B}{\Gamma \vdash \text{snd } t : B[\text{fst } t/x]}$	$\frac{\Gamma \vdash t : \cup}{\Gamma \vdash \text{El}(t)}$	$\frac{\Gamma \vdash T}{\Gamma \vdash c(T) : \cup}$	$\frac{\Gamma \vdash t : \cup}{\Gamma \vdash c(\text{El}(t)) \equiv t : \cup}$
(TYEQ/EL)	(TY/S)	(TM/s)	(TMEQ/s/ETA)	
$\frac{\Gamma \vdash T}{\Gamma \vdash \text{El}(c(T)) \equiv T}$	$\frac{\Gamma \vdash t : T}{\Gamma \vdash \mathbb{S}(t)}$	$\frac{\Gamma \vdash t : T}{\Gamma \vdash t : \mathbb{S}(t)}$	$\frac{\Gamma \vdash t_1 : T \quad \Gamma \vdash t_2 : \mathbb{S}(t_1)}{\Gamma \vdash t_1 \equiv t_2 : T}$	

Figure 6. Syntax and selected typing rules of MLTT, named variables and meta-level substitution

Brief review of MLTT. Figure 6 presents the syntax and selected typing rules of MLTT (and Figure 7 FMLTT). Dependent function types $\Pi(x : A).B$, dependent pair types $\Sigma(x : A).B$, and identity types $\text{Eq}(t_1, t_2)$ are standard. So are their introduction and elimination forms. We use

Types	$A, B, T ::= \dots \mid w\pi_1^i(\tau) \mid w\pi_2^i(\tau) \mid \mathbb{L}(\sigma) \mid \mathbb{P}(\sigma) \mid v\pi_2(\sigma) \mid \text{CaseTy}(A, B, T)$
Terms	$t, s, \ell ::= \dots \mid W(\tau) \mid W\text{sup}_i(\tau, t_1, x.t_2) \mid \mu^\bullet \mid \mu^+(\ell, x.t) \mid \text{inh}(h, \ell) \mid W\text{rec}(\tau, \ell, t) \mid \mu\pi_1(\ell) \mid \mu\pi_2(\ell) \mid v\pi_s(\sigma) \mid \mathbb{P}(\ell) \mid R\pi^i(\ell)$
W-type signatures	$\tau ::= w^\bullet \mid w^+(\tau, A, B) \mid w^-(\tau)$
Linkage signatures	$\sigma ::= v^\bullet \mid v^+(\sigma, x.s, \text{self}.T) \mid v\pi_1(\sigma) \mid \text{RecSig}(\tau, T)$
Linkage transformers	$h ::= \text{Identity} \mid \text{Extend}(h, \text{self}.t) \mid \text{Override}(h, \text{self}.t) \mid \text{Inherit}(h) \mid \text{Nest}(h, h')$

$\Gamma \vdash T_1 \equiv T_2$	$\Gamma \vdash t : T$	$\Gamma \vdash t_1 \equiv t_2 : T$	$\Gamma \vdash \tau \text{WSig}^n$	$\Gamma \vdash \sigma \text{LSig}^n$	$\Gamma \vdash h : \sigma_1 \twoheadrightarrow \sigma_2$
		(WSIG/ADD)	(TM/WSUP)		
(TM/W)	(WSIG/EMPTY)	$\Gamma \vdash \tau \text{WSig}^n$	$\Gamma \vdash \tau \text{WSig}^n$	$\Gamma \vdash t_1 : w\pi_1^i(\tau)$	
$\Gamma \vdash \tau \text{WSig}^n$	$\Gamma \vdash w^\bullet \text{WSig}^0$	$\Gamma \vdash A \quad \Gamma \vdash B : A \rightarrow \mathbb{U}$	$\Gamma, x : \text{El}(\text{app}(w\pi_2^i(\tau), t_1)) \vdash t_2 : \text{El}(W(\tau))$		
$\Gamma \vdash W(\tau) : \mathbb{U}$		$\Gamma \vdash w^+(\tau, A, B) \text{WSig}^{n+1}$	$\Gamma \vdash W\text{sup}_i(\tau, t_1, x.t_2) : \text{El}(W(\tau))$		
(TM/WREC)		(TYEQ/CASETY)			
$\Gamma \vdash \ell : \mathbb{L}(\text{RecSig}(\tau, T))$	$\Gamma \vdash t : \text{El}(W(\tau))$	$\Gamma \vdash A \quad \Gamma \vdash B : A \rightarrow \mathbb{U} \quad \Gamma \vdash T$			
$\Gamma \vdash W\text{rec}(\tau, \ell, t) : T$		$\Gamma \vdash \text{CaseTy}(A, B, T) \equiv \Pi(x : A).(\text{El}(\text{app}(B, x)) \rightarrow T) \rightarrow T$			
	(LSIG/ADD)		(L/ADD)		
(LSIG/EMPTY)	$\Gamma \vdash \sigma \text{LSig}^n$	$\Gamma, \text{self} : A \vdash T$	$\Gamma \vdash \ell : \mathbb{L}(\sigma)$	$\Gamma, \text{self} : A \vdash t : T$	
$\Gamma \vdash v^\bullet \text{LSig}^0$	$\Gamma, x : \mathbb{P}(\sigma) \vdash s : A$	$\Gamma \vdash \mu^\bullet : \mathbb{L}(v^\bullet)$	$\Gamma, x : \mathbb{P}(\sigma) \vdash s : A$		
	$\Gamma \vdash v^+(\sigma, x.s, \text{self}.T) \text{LSig}^{n+1}$		$\Gamma \vdash \mu^+(\ell, \text{self}.t) : \mathbb{L}(v^+(\sigma, x.s, \text{self}.T))$		
	(TYEQ/PK/ADD)				
	$\Gamma \vdash \sigma \text{LSig}^n$	$\Gamma, \text{self} : A \vdash T$	$\Gamma, x : \mathbb{P}(\sigma) \vdash s : A$		
	$\Gamma \vdash \mathbb{P}(v^+(\sigma, x.s, \text{self}.T)) \equiv \Sigma(x : \mathbb{P}(\sigma)).T[s/\text{self}]$				
(TMEQ/PK/ADD)			(TM/INH)		
$\Gamma \vdash \ell : \mathbb{L}(\sigma)$	$\Gamma, \text{self} : A \vdash t : T$	$\Gamma, x : \mathbb{P}(\sigma) \vdash s : A$	$\Gamma \vdash h : \sigma_1 \twoheadrightarrow \sigma_2$	$\Gamma \vdash \ell : \mathbb{L}(\sigma_1)$	
$\Gamma \vdash \mathbb{P}(\mu^+(\ell, \text{self}.t)) \equiv (\mathbb{P}(\ell), t[s[\mathbb{P}(\ell)/x]/\text{self}]) : \mathbb{P}(v^+(\sigma, s, T))$			$\Gamma \vdash \text{inh}(h, \ell) : \mathbb{L}(\sigma_2)$		
(TMEQ/OV/BETA)					
	$\Gamma \vdash h : \sigma_1 \twoheadrightarrow \sigma_2$	$\Gamma \vdash \ell : \mathbb{L}(\sigma_1)$			
$\Gamma, \text{self}_1 : A_1 \vdash t_1 : T_1$	$\Gamma, x_1 : \mathbb{P}(\sigma_1) \vdash s_1 : A_1$	$\Gamma, \text{self}_2 : A_2 \vdash t_2 : T_2$	$\Gamma, x_2 : \mathbb{P}(\sigma_2) \vdash s_2 : A_2$		
$\Gamma \vdash \text{inh}(\text{Override}(h, \text{self}_2.t_2), \mu^+(\ell, \text{self}_1.t_1)) \equiv \mu^+(\text{inh}(h, \ell), \text{self}_2.t_2) : \mathbb{L}(v^+(\sigma_2, x_2.s_2, \text{self}_2.T_2))$					

Figure 7. Syntax and selected typing rules of FMLTT, named variables and meta-level substitution

based path induction $J(\cdot, t)$ as the elimination principle for a term t of an identity type [UFP 2013]. Capture-avoiding substitution is notated $\bullet[\bullet/x]$. We use x and self to denote variables. A singleton type $\mathbb{S}(t)$ helps expose the definition of a term t in its type (rule **TM/s**) [Aspinall 1995; Stone 2000]. Definitional equalities have the forms $\Gamma_1 \equiv \Gamma_2 \vdash$, $\Gamma \vdash T_1 \equiv T_2$, $\Gamma \vdash t_1 \equiv t_2 : T$, etc. Following Altenkirch and Kaposi [2016], we regard our (intrinsically typed) syntax as being quotiented by these equalities. Quotienting facilitates coercion along equalities—given $\Gamma_1 \equiv \Gamma_2 \vdash$ and $\Gamma_1 \vdash T_1 \equiv T_2$, the derivation of $\Gamma_1 \vdash t : T_1$ is considered definitionally equal to a derivation of $\Gamma_2 \vdash t : T_2$.

We use universes à la Coquand [2013], following Kaposi et al. [2019]. Unlike Russell-style ones, these universes are not inhabited by types directly, but rather by the *codes* of types, and arguably behave better due to its closeness to Tarski-style universes [Luo 2012]. The term $\text{c}(T)$ encodes type T (**TM/c**) and thus inhabits universe \mathbb{U} . The type $\text{El}(t)$ decodes term t (**TY/El**). There is an infinite hierarchy of universes; we omit universe levels in the presentation to avoid clutter.

For concision, typing rules omit obvious premises required for well-formedness. For example, **TM/LAM** implicitly requires $\Gamma \vdash$ (i.e., that the context be well-typed).

Introducing and eliminating inductive types. *W*-types [Martin-Löf 1984] are a succinct way to model inductive types in MLTT. Together with the identity type, they can express a whole host of inductive types [Hugunin 2020], including those with multiple constructors.

Our formulation of *W*-types differs from previous ones in that it is straightforward to identify constructors from what we call *W*-type signatures. A signature $\Gamma \vdash \tau$ **WSigⁿ** is composed of n pairs of types (**WSIG/EMPTY** and **WSIG/ADD**), each modeling a constructor of the inductive type. The i -th pair, projected from the signature τ using the forms $\Gamma \vdash w\pi_1^i(\tau)$ and $\Gamma \vdash w\pi_2^i(\tau) : w\pi_1^i(\tau) \rightarrow \cup$, defines the i -th constructor. That is, given two arguments $\Gamma \vdash t_1 : w\pi_1^i(\tau)$ and $\Gamma, x : \text{El}(w\pi_2^i(\tau)(t_1)) \vdash t_2 : \text{El}(W(\tau))$, one can construct a term $\text{Wsup}_i(\tau, t_1, x.t_2)$ of the *W*-type $\text{El}(W(\tau))$ (**TM/WSUP**). $W(\tau)$ gives the code of the *W*-type (**TM/W**).

For each pair of types identifying a constructor, the first type models the non-inductive arguments of the constructor, and the second type models the *arity* of the inductive arguments. For example, the signature τ_{tm} of the *W*-type modeling the inductive type **tm** (Figure 2) is constructed as follows, where $\mathbb{0}$ is the bottom type \perp , $\mathbb{1}$ the unit type \top , $\mathbb{2}$ the boolean type, and T_{id} a type encoding id:

$$\begin{array}{llll} \text{unit} : \text{tm} & \text{var} : \text{id} \rightarrow \text{tm} & \text{abs} : \text{id} \rightarrow \text{tm} \rightarrow \text{tm} & \text{app} : \text{tm} \rightarrow \text{tm} \rightarrow \text{tm} \\ \tau_{\text{tm}}^0 := w^+(\mathbf{w}^\bullet, \mathbb{1}, \lambda_{\cdot}.\mathbb{0}) & \tau_{\text{tm}}^1 := w^+(\tau_{\text{tm}}^0, T_{\text{id}}, \lambda_{\cdot}.\mathbb{0}) & \tau_{\text{tm}}^2 := w^+(\tau_{\text{tm}}^1, T_{\text{id}}, \lambda_{\cdot}.\mathbb{1}) & \tau_{\text{tm}} := w^+(\tau_{\text{tm}}^2, \mathbb{1}, \lambda_{\cdot}.\mathbb{2}) \end{array}$$

While tm_{unit} and tm_{var} have no inductive arguments, tm_{abs} has one and tm_{app} has two. The encoding of tm_{abs} has type $T_{\text{id}} \rightarrow (\mathbb{1} \rightarrow \text{El}(W(\tau_{\text{tm}}))) \rightarrow \text{El}(W(\tau_{\text{tm}}))$, and that of tm_{app} has type $\mathbb{1} \rightarrow (\mathbb{2} \rightarrow \text{El}(W(\tau_{\text{tm}}))) \rightarrow \text{El}(W(\tau_{\text{tm}}))$. These types are strictly positive by construction.

W-types are eliminated with the form $W\text{rec}(\tau, \ell, t)$, where t is of a *W*-type $\text{El}(W(\tau))$, and ℓ is essentially an n -tuple of case handlers for the n constructors in τ (**TM/WREC**). Each case handler has a type of the form $\text{CaseTy}(A, B, T)$, where T is the motive of the recursion (**TYEQ/CASETY**); for simplicity, we model only non-dependent motives. The collection of case handlers ℓ encodes those defined and inherited by an **FRecursion** command in our plugin. We choose to type it with a linkage type $\mathbb{L}(\text{RecSig}(\tau, T))$ to avoid introducing non-dependent n -tuples, which linkages generalize.

The rules **TM/WSUP** and **TM/WREC** require access to τ : the *W*-type is exhaustively generated by its constructors, and its elimination must exhaustively handle all the constructors in its signature.

In contrast, τ should be hidden from the typing context of any term that does not invoke $\text{Wsup}(\tau, \cdot, \cdot)$ or $\text{Wrec}(\tau, \cdot, \cdot)$, so that the term can be reused—without being rechecked—for a different *W*-type signature τ' that extends τ with additional constructors. Moreover, the typing of the term should be made parametric to the definitions of those fields that invoke Wsup or Wrec , so that the term can be reused—without being rechecked—when those fields are overridden to support the extended signature τ' . Such abstraction required by family polymorphism is supported via linkages, which we discuss next.

Family polymorphism via linkages. Family polymorphism requires late binding. In FMLTT, families are expressed through linkages, and late binding of field references is achieved by requiring that typing be polymorphic to the definition of that field.

Linkage signatures have the judgment form $\Gamma \vdash \sigma$ **LSigⁿ**. A linkage signature is a list of n types (**LSIG/EMPTY** and **LSIG/ADD**). *Linkages* have the judgment form $\Gamma \vdash \ell : \mathbb{L}(\sigma)$, where $\mathbb{L}(\sigma)$ is a type formed by σ . A linkage is a list of n terms, each representing a field of the family modeled by the linkage (**L/EMPTY** and **L/ADD**).

In the rules **L/ADD** and **LSIG/ADD**, the second premise $\Gamma, \text{self} : A \vdash t : T$ is responsible for late binding. Here, A abstracts the context of the current field t , controlling how the types of the fields prior to t are exposed to the typing of t . As discussed later, the third premise $\Gamma, x : \mathbb{P}(\sigma) \vdash s : A$

$x : \mathbb{P}(\sigma_{i-1}) \vdash s_i : A_i$ $self : A_i \vdash t_i : T_i$ shown below for $i \in \{1, 2, 6, 10\}$	$\sigma_i := \nu^+(\sigma_{i-1}, x.s_i, self.T_i)$ $\ell_i := \mu^+(\ell_{i-1}, self.t_i)$ $\cdot \vdash \ell_i : \mathbb{L}(\sigma_i)$	$\mathbb{P}(\sigma_i) \equiv \Sigma(x : \mathbb{P}(\sigma_{i-1})). T_i[s_i/self]$ $\mathbb{P}(\ell_i) \equiv (\mathbb{P}(\ell_{i-1}), t_i[s_i[\mathbb{P}(\ell_{i-1})/x]/self])$ $\cdot \vdash \mathbb{P}(\ell_i) : \mathbb{P}(\sigma_i)$
Family STLC.	$\sigma_0 := \nu^\bullet$	$\ell_0 := \mu^\bullet$
FInductive $tm :=$	$x : [] \vdash s_1 : []$ $self : [] \vdash W(\tau_{tm}) : \mathbb{S}(W(\tau_{tm}))$ $\sigma_1 := \nu^+(\sigma_0, x.s_1, self.T_1)$	$\ell_1 := \mu^+(\ell_0, self.t_1)$
$tm_unit : tm$	$x : [tm : \mathbb{S}(W(\tau_{tm}))] \vdash s_2 : [tm : \mathbb{S}(W(\tau_{tm}))]$ $self : [tm : \mathbb{S}(W(\tau_{tm}))] \vdash Wsup(\tau_{tm}, \top, x.\perp) : El(self \triangleright tm)$ $\sigma_2 := \nu^+(\sigma_1, x.s_2, self.T_2)$	$\ell_2 := \mu^+(\ell_1, self.t_2)$
...
FRecursion subst ...	<i>$W(\tau_{tm})$ is abstracted into \cup in the typing context (i.e., A_6) of the case handler below:</i>	
Case $tm_unit := \dots$	$x : [tm : \mathbb{S}(W(\tau_{tm})), tm_unit : El(tm), \dots] \vdash s_6 : [tm : \cup, tm_unit : El(tm), \dots]$ $self : [tm : \cup, tm_unit : El(tm), \dots] \vdash \dots : CaseTy\left(\top, \perp, \begin{matrix} El(self \triangleright tm) \rightarrow \\ id \rightarrow El(self \triangleright tm) \end{matrix}\right)$ $\sigma_6 := \nu^+(\sigma_5, x.s_6, self.T_6)$	$\ell_6 := \mu^+(\ell_5, self.t_6)$
Case ... Case ... Case
End subst.	$x : \left[\begin{matrix} tm : \mathbb{S}(W(\tau_{tm})), tm_unit : El(tm), \dots \\ subst_tm_unit : CaseTy(\top, \perp, \dots), \dots \end{matrix} \right] \vdash s_{10} : \left[\begin{matrix} tm : \mathbb{S}(W(\tau_{tm})), tm_unit : El(tm), \dots \\ subst_tm_unit : CaseTy(\top, \perp, \dots), \dots \end{matrix} \right]$ $self : \left[\begin{matrix} tm : \mathbb{S}(W(\tau_{tm})), tm_unit : El(tm), \dots \\ subst_tm_unit : CaseTy(\top, \perp, \dots), \dots \end{matrix} \right] \vdash \lambda x. Wrec(\tau_{tm}, \dots, x) : \begin{matrix} El(self \triangleright tm) \rightarrow \\ El(self \triangleright tm) \rightarrow \\ id \rightarrow El(self \triangleright tm) \end{matrix}$ $\sigma_{10} := \nu^+(\sigma_9, x.s_{10}, self.T_{10})$	$\ell_{10} := \mu^+(\ell_9, self.t_{10})$
...
End STLC.	$\cdot \vdash \mathbb{P}(\ell_{53}) : \mathbb{P}(\sigma_{53})$	

Figure 8. FMLTT encoding of the STLC family from Figure 2. Each row $self : A_i \vdash t_i : T_i$ types a field. The type A_i controls how the typing of field t_i sees the types of the fields prior to t_i .

is responsible for creating the context type A that possibly hides W -type signatures in σ , which records the types of the prior fields.

Crucially, the premise $\Gamma, self : A \vdash t : T$ makes clear that the typing of t in **L/ADD** is *universally* quantified—rather than *existentially* quantified as is in **TM/PAIR**—over how the fields in t 's context are defined. Late binding enables reuse. A different linkage ℓ' that overrides fields in t 's context (and thus models a derived family) can reuse t —without retyping it—by first projecting t from $\mu^+(\ell, self.t)$ and then appending it to ℓ' .

The type A in **L/ADD** and **LSIG/ADD**, abstracting the types of the prior fields, does not necessarily contain the same types as those recorded by $\mathbb{L}(\sigma)$, because a field defined as the code $W(\tau)$ of a W -type has to expose different types to different fields that come after it. Later fields that invoke $Wsup(\tau, \cdot, \cdot)$ or $Wrec(\tau, \cdot, \cdot)$ should see the concrete signature τ , as the rules **TM/WSUP** and **TM/WREC** stipulate. By contrast, τ should be hidden from all other fields, so that they can be reused in a different context where the W -type signature τ is replaced by an extended one τ' .

We use Figure 8 to illustrate. On the left is code excerpted from Figure 2, and on the right is how the corresponding fields are modeled and typed in FMLTT. ℓ_i is the linkage that adds the i -th field t_i with the typing $\mathit{self} : A_i \vdash t_i : T_i$. The context types A_i are a dependent tuple type, but for readability, we write them as dependent record types that give labels to the fields. Field accesses are notated, for example, as $\mathit{self}_{\triangleright} \mathit{tm}$.

- The first field t_1 is defined as $W(\tau_{\mathit{tm}})$ and given the singleton type $\mathbb{S}(W(\tau_{\mathit{tm}}))$, where τ_{tm} is the W -type signature constructed earlier for the extensible inductive type tm . This typing is recorded by σ_1 and is thus available in all σ_i 's.
- The next four fields model the four constructors of tm . Constructor $\mathit{tm_unit}$ is modeled as $W\mathit{sup}(\tau_{\mathit{tm}}, \top, x, \perp)$ and has type $\mathbb{E}l(\mathit{self}_{\triangleright} \mathit{tm})$, where self stands for the typing context containing the first field tm . The W -type signature τ_{tm} is exposed in this typing context; that is, $\mathit{self}_{\triangleright} \mathit{tm}$ has type $\mathbb{S}(W(\tau_{\mathit{tm}}))$. So $\mathbb{E}l(\mathit{self}_{\triangleright} \mathit{tm})$ and $\mathbb{E}l(W(\tau_{\mathit{tm}}))$ can be equated, as required by rule **TM/WSUP**.
- Likewise, τ_{tm} is exposed in the typing context of t_{10} , which models the recursive function subst by invoking the recursor $W\mathit{rec}(\tau_{\mathit{tm}}, \ast, \ast)$. Like subst , partial recursors in **FPOP** are axiomatized by the plugin (Sections 3.2 and 3.6), and they can similarly be defined in FMLTT using $W\mathit{rec}$.
- By contrast, τ_{tm} is hidden from the typing of all other fields. Their typing should depend on the knowledge that tm has type \mathbb{U} , rather than $\mathbb{S}(W(\tau_{\mathit{tm}}))$, so that they can be reused in a context where tm is defined as $W(\tau'_{\mathit{tm}})$, where τ'_{tm} extends τ_{tm} with additional constructors. For example, in Figure 8, the typing of the case handlers of subst (e.g., t_6) is oblivious to the definition of tm —it sees only $\mathit{tm} : \mathbb{U}$ —so the case handlers can be reused by a linkage modeling **STLCFix**.

In **L/ADD** and **LSIG/ADD**, the third premise $\Gamma, x : \mathbb{P}(\sigma) \vdash s : A$ is responsible for hiding W -type signatures. Here, $\mathbb{P}(\sigma)$ packages σ —recall that σ contains the (self-parameterized) types of all the fields preceding the current field t —into a dependent tuple type (**TYPEQ/PK/ADD**). The term s turns a tuple of type $\mathbb{P}(\sigma)$ into a new tuple of type A that hides W -type signatures behind \mathbb{U} , if necessary.⁴

It is straightforward to find the s that fits the bill, though this process is not automated in FMLTT. In particular, when no hiding is needed (that is, when the field being checked invokes a W -type constructor or eliminator), s is simply x , and A is $\mathbb{P}(\sigma)$. In Figure 8, s_2 and s_{10} are x . Otherwise, it is needed to hide W -type signatures. In Figure 8, s_6 hides $\mathit{tm} : \mathbb{S}(W(\tau_{\mathit{tm}}))$ as $\mathit{tm} : \mathbb{U}$ in A_6 , so that t_6 , typed under $\mathit{self} : A_6$, is oblivious to the concrete signature τ_{tm} and therefore can be reused.

When a family is concluded (e.g., **End STLC**), a linkage ℓ containing all the fields is available (e.g., ℓ_{53} in Figure 8). Fields of the family can then be accessed by projecting them out of the tuple $\mathbb{P}(\ell)$. As **TMEQ/PK/ADD** indicates, $\mathbb{P}(\ell)$ ties the recursive knot: it packages the linkage ℓ into a dependent tuple of type $\mathbb{P}(\sigma)$, by instantiating the self parameters.

Linkage transformers. Inheritance and code reuse can already be expressed through the projection of fields out of linkages and their inclusion into new linkages. To make common patterns of linkage manipulations more convenient, FMLTT provides a “library” of *linkage transformers*, whose well-formedness judgments have form $\Gamma \vdash h : \sigma_1 \twoheadrightarrow \sigma_2$. The idea is that applying h to a linkage of type $\mathbb{L}(\sigma_1)$ yields a linkage of type $\mathbb{L}(\sigma_2)$ (**TM/INH**).

Derived families can be modeled as linkage transformers inductively constructed from the introduction forms **Identity**, **Extend**($h, \mathit{self}.t$), **Override**($h, \mathit{self}.t$), etc. Figure 7 shows the β -rule of an example transformer, **TMEQ/OV/BETA**. It states that applying the transformer **Override**($h, \mathit{self}_2.t_2$) to a linkage of form $\mu^+(\ell, \mathit{self}_1.t_1)$ overrides the linkage’s last field t_1 with t_2 . For instance, the construction below shows that **Override**(**Identity**, $\mathit{self}.W(\tau'_{\mathit{tm}})$) is used as the first step in creating a linkage transformer modeling a derived family that overrides τ_{tm} with an extended signature τ'_{tm} :

⁴In the presence of hiding, it is important that the context type A be a dependent tuple type rather than a linkage type; we expand on this point in Section 6.

Family STLCFix extends STLC. $h_0 := \text{Identity}$

$\cdot \vdash h_0 : v^\bullet \Rightarrow v^\bullet$

FInductive tm += ...

$h_1 := \text{Override}(h_0, \text{self}.W(\tau'_{\text{tm}})) \cdot \vdash h_1 : \frac{v^+(v^\bullet, x.s_1, \text{self}.S(W(\tau_{\text{tm}}))) \rightarrow v^+(v^\bullet, x'.s'_1, \text{self}'.S(W(\tau'_{\text{tm}})))}$

Section 6.5 sketches how the other introduction forms of linkage transformers can be used to model the construction of a derived family as a linkage transformer.

Why $\mathbb{P}(\sigma)$? $\mathbb{P}(\sigma)$ is needed in rules **LSIG/ADD** and **L/ADD**: these rules have a premise $\Gamma, x : \mathbb{P}(\sigma) \vdash s : A$ responsible for possibly hiding W-type signatures (in $\mathbb{P}(\sigma)$) behind \cup (in A); see s_6 in Figure 8 for an example. Here, both $\mathbb{P}(\sigma)$ and A are dependent tuple types rather than linkage types.

An alternative to $\mathbb{P}(\sigma)$ would be to use linkage types $\mathbb{L}(\cdot)$ in this premise. However, this alternative is infeasible. We illustrate using a simple example. Consider the linkage signature σ below (it is written using the record syntax for illustration but should not be confused with records, as a later field is universally quantified over a *self* variable). The job is to hide the W-types $\mathbb{S}(W(\tau_1))$ and $\mathbb{S}(W(\tau_2))$ in σ behind \cup in σ^A :

$$\begin{aligned} \sigma &:= \{T_A : \mathbb{S}(W(\tau_1)); & T_B : \mathbb{S}(W(\tau_2)); & f : \text{El}(\text{self}_\triangleright T_A) \rightarrow \text{El}(\text{self}_\triangleright T_B)\} \\ \sigma^A &:= \{T_A : \cup; & T_B : \cup; & f : \text{El}(\text{self}_\triangleright T_A) \rightarrow \text{El}(\text{self}_\triangleright T_B)\} \\ & & & x : \mathbb{L}(\sigma) \vdash s : \mathbb{L}(\sigma^A) \end{aligned}$$

The problem lies in that the type $\mathbb{L}(\sigma^A)$ above cannot be inhabited. Suppose there is some ℓ^A such that $\cdot \vdash \ell^A : \mathbb{L}(\sigma^A)$. Then it follows from the self-parameterization that projecting out the last field f from ℓ^A (notated as $\ell^A.f$ for illustration) has the typing $T_A : \cup, T_B : \cup \vdash \ell^A.f : \text{El}(T_A) \rightarrow \text{El}(T_B)$. This typing means that there is a function of type $\Pi(T_A : \cup).\Pi(T_B : \cup).\text{El}(T_A) \rightarrow \text{El}(T_B)$, i.e., a function from an arbitrary type to another arbitrary type. We can easily derive inconsistency from this function: applying it to \top and \perp yields a function of type $\top \rightarrow \perp$.

Thus, it is critical that the hiding of W-type signatures operate on dependent tuple types rather than linkage types. To this end, we introduce the syntax $\mathbb{P}(\sigma)$ and its computation rules (e.g., **TYEQ/PK/ADD**) to make it convenient to package linkage signatures into dependent tuple types.

The complete formalization. The definitive version containing all the rules in FMLTT is stated in Section 6. This definitive version uses de Bruijn indices and explicit substitution. An effort to mechanize the formalization is underway. It follows [Altenkirch and Kaposi \[2016\]](#) in using Quotient Inductive-Inductive Type (**QIIT**) as the logical framework to state intrinsically typed syntax (and also in using Agda's **REWRITE** pragma to replace the limited uses of equality reflection, as Agda does not support QIITs natively).

Consistency and canonicity. One of the most fundamental properties of a dependent type theory is consistency.

THEOREM 5.1 (CONSISTENCY). *The typing judgment $\cdot \vdash t : \perp$ is not derivable for any term t .*

Consistency says that the type \perp is not inhabited. Thus, it is safe to use the type theory for logical reasoning, as not every proposition is trivially provable.

A second property we prove is canonicity.

THEOREM 5.2 (CANONICITY). *If $\cdot \vdash t : \mathbb{B}$, then either $\cdot \vdash t \equiv \text{tt} : \mathbb{B}$ or $\cdot \vdash t \equiv \text{ff} : \mathbb{B}$.*

This canonicity theorem says that every closed term of the ground type \mathbb{B} is convertible to one of the canonical forms **tt** and **ff**. When the canonicity theorem is proved in a constructive metalogic, its proof amounts to a normalization function for closed terms of the ground type. So canonicity serves to justify the computational nature of the type theory.

We prove Theorems 5.1 and 5.2 by constructing a logical-relations model for the well-formedness rules, following prior approaches [Coquand 2019; Kaposi et al. 2019; Sterling 2019]. The model interprets a linkage $\mu^+(\ell, t)$ as a pair where the second component is a function (modeling late binding), as rule **L/ADD** indicates. The model interprets the bottom type \perp as an empty set, from which Theorem 5.1 follows. A closed, well-formed type $\cdot \vdash T$ is interpreted as a logical predicate on closed terms: the predicate includes all closed, “reducible” terms of type T . Theorem 5.2 follows from this interpretation. The detailed construction of the logical-relations model is available in Section 6.4.

6 SYNTAX AND SEMANTIC MODELS OF FMLTT

This section supplements Section 5 with the details of the syntax and the proofs of the main results.

Unlike in Section 5, which uses named binders, here we use *de Bruijn* indices and *explicit substitutions* [Abadi et al. 1989; Martin-Löf 1992]: substitutions γ and their applications (e.g., $T[\gamma]$, which applies γ to the type T) are part of the syntax rather than meta-operations. We work in an intrinsically typed setting: terms are well typed by construction. Consequently, we omit without ambiguity some obvious premises needed for well-formedness. This style of syntax formulation follows a recent trend [Altenkirch and Kaposi 2016; Coquand 2019; Gratzer et al. 2019] known as the “algebraic presentation” of MLTT. Moreover, universe levels are explicit here. Other aspects of the syntax remain the same as in Section 5.

We still use based path induction as the elimination principle for identity types [Univalent Foundations Program 2013].

6.1 MLTT with Explicit Substitutions and Universe Levels

We review the base MLTT fragment of FMLTT first. The following rules are the fully expanded version of Figure 6, where each type (except universes) has four kinds of rules: introduction, elimination, computation and uniqueness rules. Computation rules (β -equalities) indicate the judgemental equality for expressing reduction when elimination rules apply to the introduction form. Uniqueness rules (η -equalities) indicate how some terms are possibly uniquely constructed by introduction rules. The substitution rule is in the form of de Bruijn indices.

The first four rules are following the style of Univalent Foundations Program [2013]; Pfennig [2009], while the substitution rules are naturally required for any language with first-class functions.

$$\begin{array}{l}
\text{Contexts } \Gamma, \Delta, \Theta ::= \cdot \mid \Gamma, A \\
\text{Substitutions } \gamma ::= p^n \mid \gamma, t \mid \gamma_1 \circ \gamma_2 \mid \pi_1 \gamma \mid \text{id} \\
\text{Types } A, B, T ::= T[\gamma] \mid \mathbb{U} \mid \mathbb{B} \mid \perp \mid \top \mid \Pi(A, B) \mid \Sigma(A, B) \mid \text{Eq}(t_1, t_2) \mid \mathbb{S}(t) \mid \text{El}(t) \\
\text{Terms } t, s ::= t[\gamma] \mid \text{var}_n \mid \pi_2 \gamma \mid \text{c}(T) \mid () \mid \text{tt} \mid \text{ff} \mid \text{if}(t_1, t_2, t_3) \mid \lambda(t) \mid \text{app}(t) \mid \\
\quad (t_1, t_2) \mid \text{fst } t \mid \text{snd } t \mid \text{refl}(t) \mid J(t_1, t_2)
\end{array}$$

$$\begin{array}{c}
\boxed{\Gamma \vdash_k} \qquad \boxed{\Gamma \vdash \gamma : \Delta} \qquad \boxed{\Gamma \vdash_j T} \qquad \boxed{\Gamma \vdash t : T} \\
\frac{}{\Gamma \vdash_0 \cdot} \qquad \frac{\Gamma \vdash_k \quad \Gamma \vdash_j A}{\Gamma, A \vdash_{k \sqcup j}} \\
\frac{}{\Gamma \vdash_{j+1} \mathbb{U}_j} \qquad \frac{}{\Gamma \vdash_0 \mathbb{B}} \qquad \frac{}{\Gamma \vdash_0 \perp} \qquad \frac{}{\Gamma \vdash_0 \top} \qquad \frac{\Gamma \vdash_j A \quad \Gamma, A \vdash_i B}{\Gamma \vdash_{j \sqcup i} \Pi(A, B)} \qquad \frac{\Gamma \vdash_j A \quad \Gamma, A \vdash_i B}{\Gamma \vdash_{j \sqcup i} \Sigma(A, B)} \\
\frac{\Gamma \vdash_j A \quad \Gamma \vdash x : A \quad \Gamma \vdash y : A}{\Gamma \vdash_j \text{Eq}(x, y)} \qquad \frac{\Gamma \vdash_j A \quad \Gamma \vdash a : A}{\Gamma \vdash_j \mathbb{S}(a)} \qquad \frac{\text{(TY/SUB)} \quad \Delta \vdash_j T \quad \Gamma \vdash \gamma : \Delta}{\Gamma \vdash_j T[\gamma]} \\
\frac{}{\Gamma \vdash A[p^0] \equiv A} \qquad \frac{}{\Gamma \vdash A[\gamma_1 \circ \gamma_2] \equiv A[\gamma_1][\gamma_2]} \\
\frac{}{\Gamma \vdash \gamma : \Delta} \\
\frac{\Gamma \vdash \mathbb{U}[\gamma] \equiv \mathbb{U} \quad \Gamma \vdash \mathbb{B}[\gamma] \equiv \mathbb{B} \quad \Gamma \vdash \perp[\gamma] \equiv \perp \quad \Gamma \vdash (\Pi(A, B))[\gamma] \equiv \Pi(A[\gamma], B[\gamma^\uparrow])}{\Gamma \vdash (\Sigma(A, B))[\gamma] \equiv \Sigma(A[\gamma], B[\gamma^\uparrow])} \quad \Gamma \vdash (\text{Eq}(a, b))[\gamma] \equiv \text{Eq}(a[\gamma], b[\gamma]) \\
\Gamma \vdash \mathbb{S}(a)[\gamma] \equiv \mathbb{S}(a[\gamma])
\end{array}$$

$$\begin{array}{c}
\frac{\Gamma \vdash_j T}{\Gamma \vdash c(T) : \mathbb{U}_j} \quad \frac{\Gamma \vdash T : \mathbb{U}_j}{\Gamma \vdash_j \text{El}(T)} \quad \frac{}{\Gamma \vdash \text{El}(c(T)) \equiv T} \quad \frac{}{\Gamma \vdash c(\text{El}(T)) \equiv T : \mathbb{U}} \quad \frac{}{\Gamma \vdash () : \mathbb{T}} \\
\\
\frac{\Gamma \vdash t : \mathbb{T}}{\Gamma \vdash t \equiv () : \mathbb{T}} \quad \frac{\text{(TM/SUB)} \quad \frac{\Delta \vdash t : T \quad \Gamma \vdash \gamma : \Delta}{\Gamma \vdash t[\gamma] : T[\gamma]}}{\Gamma \vdash t[\mathfrak{p}^0] \equiv t : T} \quad \frac{\Gamma, A \vdash t : B}{\Gamma \vdash \lambda(t) : \Pi(A, B)} \\
\\
\frac{\Gamma \vdash t : \Pi(A, B)}{\Gamma, A \vdash \text{app}(t) : B} \quad \frac{}{\Gamma, A \vdash \text{app}(\lambda(t)) \equiv t : B} \quad \frac{}{\Gamma \vdash \lambda(\text{app}(t)) \equiv t : \Pi(A, B)} \\
\\
\frac{\Gamma \vdash u : A \quad \Gamma \vdash v : B[(\mathfrak{p}^0, u)]}{\Gamma \vdash (u, v) : \Sigma(A, B)} \quad \frac{}{\Gamma \vdash \text{fst } t : A \quad \Gamma \vdash \text{snd } t : B[(\mathfrak{p}^0, \text{fst } t)]} \\
\Gamma \vdash \text{fst } (u, v) \equiv u : A \quad \Gamma \vdash \text{snd } (u, v) \equiv v : B[(\mathfrak{p}^0, u)] \quad \Gamma \vdash (\text{fst } t, \text{snd } t) \equiv t : \Sigma(A, B) \\
\\
\frac{}{\Gamma \vdash \text{tt}, \text{ff} : \mathbb{B}} \quad \frac{\Gamma \vdash c : \mathbb{B} \quad \Gamma \vdash a : T \quad \Gamma \vdash b : T}{\Gamma \vdash \text{if}(c, a, b) : T} \quad \frac{\Gamma \vdash a : A}{\Gamma \vdash \text{refl}(a) : \text{Eq}(a, a)} \\
\frac{\Gamma \vdash u : A \quad \Gamma, A, \text{Eq}(u[\pi_1], \pi_2) \vdash C \quad \Gamma \vdash w : C[\mathfrak{p}^0, u, \text{refl}(u)] \quad \Gamma \vdash v : A \quad \Gamma \vdash t : \text{Eq}(u, v)}{\Gamma \vdash \text{J}(w, t) : C[\mathfrak{p}^0, v, t]} \\
\Gamma \vdash \text{if}(\text{tt}, a, b) \equiv a : T \quad \Gamma \vdash \text{if}(\text{ff}, a, b) \equiv b : T \quad \Gamma \vdash \text{J}(w, \text{refl}(u)) \equiv w : C[\mathfrak{p}^0, u, \text{refl}(u)] \\
\\
\frac{\Gamma \vdash a : A}{\Gamma \vdash a : \mathbb{S}(a)} \quad \frac{\Gamma \vdash_A a \quad \Gamma \vdash x : \mathbb{S}(a)}{\Gamma \vdash x \equiv a : A} \\
\\
\Gamma \vdash (\lambda(t))[\gamma] \equiv \lambda(t[\gamma^\uparrow]) : \Pi(A, B) \quad \Gamma \vdash (u, v)[\gamma] \equiv (u[\gamma], v[\gamma]) : \Sigma(A, B) \\
\Gamma \vdash \text{El}(T[\gamma]) \equiv (\text{El}(T))[\gamma] \quad \Gamma \vdash \text{tt}[\gamma] \equiv \text{tt} : \mathbb{B} \quad \Gamma \vdash \text{ff}[\gamma] \equiv \text{ff} : \mathbb{B} \\
\Gamma \vdash (\text{if}(c, a, b))[\gamma] \equiv \text{if}(c[\gamma], a[\gamma], b[\gamma]) : T \quad \Gamma \vdash (\text{J}(w, t))[\gamma] \equiv \text{J}(w[\gamma], t[\gamma]) : \\
\\
\frac{}{\Gamma \vdash \epsilon : \cdot} \quad \frac{\Gamma \vdash \gamma : \cdot}{\Gamma \vdash \gamma \equiv \epsilon : \cdot} \quad \frac{\Delta \vdash \delta : \Theta \quad \Gamma \vdash \gamma : \Delta}{\Gamma \vdash \delta \circ \gamma : \Theta} \quad \frac{}{\Gamma \vdash \text{id} \equiv \mathfrak{p}^0 : \Gamma} \\
\\
\text{(SUB/ID)} \quad \frac{}{\Gamma \vdash \mathfrak{p}^0 : \Gamma} \quad \text{(SUB/EXT)} \quad \frac{\Gamma \vdash \gamma : \Delta \quad \Gamma \vdash t : A[\gamma]}{\Gamma \vdash \gamma, t : (\Delta, A)} \quad \text{(SUB/WK)} \quad \frac{\Gamma \vdash \mathfrak{p}^n : \Delta \quad \Gamma \vdash A}{\Gamma, A \vdash \mathfrak{p}^{n+1} : \Delta} \\
\\
\text{(TM/VAR)} \quad \frac{\Gamma, A_n, \dots, A_1, A_0 \vdash}{\Gamma, A_n, \dots, A_1, A_0 \vdash \text{var}_n : A_n[\mathfrak{p}^{n+1}]} \quad \text{(SUB/DBJ/SHIFT)} \quad \frac{\Gamma \vdash \gamma : \Delta \quad \Delta \vdash A}{\Gamma, A[\gamma] \vdash \gamma^{\uparrow A} \equiv (\gamma \circ \mathfrak{p}^1, \text{var}_0) : \Delta, A} \\
\\
\frac{\Gamma \vdash \gamma : (\Delta, A)}{\Gamma \vdash \pi_1 \gamma : \Delta} \quad \frac{\Gamma \vdash \gamma : (\Delta, A)}{\Gamma \vdash \pi_2 \gamma : A[\pi_1 \gamma]} \quad \frac{}{\Gamma \vdash (\pi_1 \gamma, \pi_2 \gamma) \equiv \gamma : \Delta} \\
\Gamma \vdash \gamma_1 \circ (\gamma_2 \circ \gamma_3) \equiv (\gamma_1 \circ \gamma_2) \circ \gamma_3 : \Theta \quad \Gamma \vdash \mathfrak{p}^0 \circ \gamma \equiv \gamma \circ \mathfrak{p}^0 \equiv \gamma : \Theta
\end{array}$$

MLTT as a programming language. Conventionally, the operational semantics sketches *the interpreter* of a given formalization of a programming language. However, compared to this traditional setup, a careful reader will notice the absence of operational semantics in our MLTT. It is no doubt that MLTT is a logical system where we can write proofs, but it is unclear if MLTT can really be considered as a programming language.

Indeed, for MLTT, we only have syntactic term and their judgemental equality (conversion rules) among them. However, we still can have *an interpreter* about our MLTT in our ambient meta-logic. For a programmer, this is like writing a Haskell meta-circular interpreter [Abelson and Sussman 1996] that interprets an abstract syntax tree into Haskell’s data structure. For a logician, this is interpreting the syntactic $\cdot \vdash t : \mathbb{B}$ to an element $\llbracket t \rrbracket \in \{\text{tt}, \text{ff}\}$. For a formal semanticist, this can be considered as a naive form of denotational semantic.

This interpreter is in the form of our canonicity theorem [Theorem 5.2](#) in the light of Curry-Howard Correspondence, as any closed program of type \mathbb{B} is evaluated to either tt or ff under this theorem. What’s more, our canonicity theorem also *reflects* this computational traces into MLTT as a proof term of the derivation $t \equiv \text{tt}$ or $t \equiv \text{ff}$. From this perspective, we can see the similarity between the unidirectional small-step relation and the bidirectional conversion rules.

Quotient upon the syntax. As we mentioned in [Section 5](#), the syntax is considered as being quotiented by the judgmental equalities, and quotienting facilitates type coercion along equalities.

By looking at our syntax as some sorts of abstract syntax tree (data), after quotient, each syntax piece is actually an equivalent class of unquotiented syntax, where the equivalence is exactly the judgemental equalities. A proper introduction of this quotient concept can be found in [Altenkirch and Kaposi \[2016\]](#); [Kaposi \[2017\]](#), where a type-theoretic quotient formulation of MLTT is given.

Once it is clear that the syntax of our MLTT as a logic system is a quotiented data, then it is natural to consider model of this logic system as a function mapping from those quotiented data (syntax). The (semantic) model introduced below will follow this idea. More specifically, the semantic model as a function will have to respect judgemental equalities to be a well-defined function from quotiented data. In other words, for two judgementally equal syntax pieces, the model needs to map them into the same data.

De Bruijn Indices and Explicit Substitution. [De Bruijn](#) indices and explicit substitutions make details about binders and substitutions explicit. Using explicit substitutions obviates the need for special treatment of substitutions in the proofs, as substitutions are part of the syntax. The form var_n represents a variable bound by the n -th closest enclosing binder. For example, $\lambda x. \lambda y. x$ is $\lambda(\lambda(\text{var}_1))$. Substitutions are typed with the form $\Gamma \vdash \gamma : \Delta$. The idea is that applying γ to terms valid in the context Δ yields terms valid in Γ ([TM/SUB](#) and [TY/SUB](#)). The two main forms of substitutions are weakening ([SUB/WK](#)) and extension ([SUB/EXT](#)): $t[p^n]$ introduces n free variables into the context of t , and $t[\gamma, t']$ substitutes t' for var_0 in t and then applies γ . For example, rule [TM/SND](#) states that if t is a dependent pair that has type $\Sigma(A, B)$, then $\text{snd } t$ has type $B[p^0, \text{fst } t]$, where p^0 is the identity substitution ([SUB/ID](#)). We occasionally use the notation id for p^0 .

To simplify, p^n is a short hand for $\pi_1^n \text{ id}$ and var_n is a short hand for $\pi_2 \pi_1^n$. Thus during meta-theoretic reasoning, we will only deal with π_1 and π_2 .

Consequently, function application changes to an equivalent formulation, and becomes a “direct inverse” of typing rule for function abstraction. For example, the named notation $\text{app}(f, t)$ can be equivalently represented by $\text{app}(f)[p^0, t]$.

Finally, we have [SUB/DBJ/SHIFT](#) defined using [SUB/WK](#) and [TM/VAR](#). This rule applying substitution γ to the earlier portion of the context. We usually omit A in the $\gamma^{\uparrow A}$ because it can be inferred from the context.

Universe levels. Universe levels address the size issue—it is unsound to have a set of all sets (or a universe of all types) [[Hurkens 1995](#)]. The level of a universe specifies “how large” that universe is.

The judgment form $\Gamma \vdash_i T$ indicates that (the code of) the type T inhabits universe \mathcal{U}_i . Similarly, $\Gamma \vdash_i$ indicates that (the codes of) the types in Γ inhabit universe \mathcal{U}_i . The notation $i \sqcup j$ denotes $\max(i, j)$.

6.2 FMLTT

FMLTT extends the MLTT in Section 6.1 with W-type signatures and W-types, linkage signatures and linkages, and linkage transformers. The following rules are the fully expanded version of Figure 7, where each type has five kinds of rules: introduction, elimination, computation, uniqueness rules and substitution lemma. Computation rules (β -equalities) indicate the judgemental equality for expressing reduction when elimination rules apply to the introduction form. Uniqueness rules (η -equalities) indicate how some terms are possibly uniquely constructed by introduction rules. The substitution rule is in the form of de Bruijn indices.

The first four rules are following the style of **Univalent Foundations Program [2013]**; **Pfenning [2009]**, while the substitution rules are naturally required for any language with first-class functions.

Apart from types, we also lay out the complete version of judgements for signatures and linkage transformers. The signature judgements are also following these five parts—introduction, elimination, computation, uniqueness and substitution rules.

Types	$A, B, T ::= \dots \mid w\pi_1^i(\tau) \mid w\pi_2^i(\tau) \mid \mathbb{L}(\sigma) \mid \mathbb{P}(\sigma) \mid v\pi_2(\sigma) \mid \text{CaseTy}(A, B, T)$
Terms	$t, s, \ell ::= \dots \mid W(\tau) \mid \text{Wsup}_i(\tau, t_1, t_2) \mid \mu^\bullet \mid \mu^+(\ell, t) \mid \text{inh}(h, \ell) \mid \text{Wrec}(\tau, \ell, t) \mid \mu\pi_1(\ell) \mid \mu\pi_2(\ell) \mid v\pi_s(\sigma) \mid \mathbb{P}(\ell) \mid R\pi^i(\ell)$
W-type signatures	$\tau ::= w^\bullet \mid w^+(\tau, A, B) \mid \tau[\gamma] \mid w^-(\tau)$
Linkage signatures	$\sigma ::= v^\bullet \mid v^+(\sigma, s, T) \mid v\pi_1(\sigma) \mid \text{RecSig}(\tau, T) \mid \sigma[\gamma]$
Linkage transformers	$h ::= \text{Identity} \mid \text{Extend}(h, t) \mid \text{Override}(h, t) \mid \text{Inherit}(h) \mid \text{Nest}(h, h') \mid h[\gamma]$

$$\begin{array}{c}
 \boxed{\Gamma \vdash_m \tau \text{WSig}^n} \qquad \boxed{\Gamma \vdash_l \sigma \text{LSig}^n} \qquad \boxed{\Gamma \vdash h : \sigma_1 \twoheadrightarrow \sigma_2} \\
 \\
 \frac{}{\Gamma \vdash_m w^\bullet \text{WSig}^0} \qquad \frac{\Gamma \vdash_m \tau \text{WSig}^0}{\Gamma \vdash_m \tau \equiv w^\bullet \text{WSig}^0} \qquad \frac{\Gamma \vdash_m \tau \text{WSig}^n \quad \Gamma \vdash_m A \quad \Gamma, A \vdash_m B}{\Gamma \vdash_m w^+(\tau, A, B) \text{WSig}^{n+1}} \\
 \\
 \frac{\Gamma \vdash_m \tau \text{WSig}^n \quad j < n}{\Gamma \vdash_m w\pi_1^j(\tau) \quad \Gamma, w\pi_1^j(\tau) \vdash_m w\pi_2^j(\tau)} \\
 \\
 \frac{\Gamma \vdash \gamma : \Theta}{\Gamma \vdash w\pi_1^{j+1}(w^+(\tau, A, B)) \equiv w\pi_1^j(\tau) \quad \Gamma, w\pi_1^{j+1}(w^+(\tau, A, B)) \vdash w\pi_2^{j+1}(w^+(\tau, A, B)) \equiv w\pi_2^j(\tau)} \\
 \frac{\Gamma \vdash w\pi_1^0(w^+(\tau, A, B)) \equiv A \quad \Gamma, w\pi_1^0(w^+(\tau, A, B)) \vdash w\pi_2^0(w^+(\tau, A, B)) \equiv B}{\Gamma \vdash_m \tau \text{WSig}^{n+1}} \\
 \frac{\Gamma \vdash_m \tau \text{WSig}^{n+1}}{\Gamma \vdash_m w^-(\tau) \text{WSig}^n} \\
 \\
 \frac{\Gamma \vdash \gamma : \Theta}{\Gamma \vdash w^\bullet[\gamma] \equiv w^\bullet \text{WSig}^0 \quad \Gamma \vdash w^+(\tau, A, B)[\gamma] \equiv w^+(\tau[\gamma], A[\gamma], B[\gamma^\uparrow]) \text{WSig}^{n+1}} \\
 \frac{\Gamma \vdash w\pi_1^j(\tau)[\gamma] \equiv w\pi_1^j(\tau[\gamma]) \quad \Gamma, w\pi_1^j(\tau[\gamma]) \vdash_i w\pi_2^j(\tau)[\gamma^\uparrow] \equiv w\pi_2^j(\tau[\gamma])}{\Gamma \vdash_m \tau \text{WSig}^n \quad \Gamma \vdash t_1 : w\pi_1^i(\tau)} \\
 \frac{\Gamma \vdash_m \tau \text{WSig}^n}{\Gamma \vdash W(\tau) : \mathbb{U}_{m+1}} \qquad \frac{\Gamma \vdash_m \tau \text{WSig}^n \quad \Gamma \vdash t_1 : w\pi_1^i(\tau)}{\Gamma \vdash \text{Wsup}_i(\tau, t_1, t_2) : \text{El}(W(\tau))}
 \end{array}$$

$$\begin{array}{c}
\Gamma \vdash \gamma : \Theta \\
\hline
\Gamma \vdash (\text{Wsup}(T, a, b))[Y] \equiv \text{Wsup}(T[Y], a[Y], b[Y^\uparrow]) : \text{El}((\text{W}(\tau[Y]))) \\
\hline
\frac{\Gamma \vdash_m \tau \text{WSig}^n \quad \Gamma \vdash_m A \quad \Gamma, A \vdash_m B}{\Gamma \vdash \text{w}^-(\text{w}^+(\tau, A, B)) \equiv \tau \text{WSig}^n} \quad \frac{\Gamma \vdash \gamma : \Theta}{\Gamma \vdash \text{w}^-(\tau)[Y] \equiv \text{w}^-(\tau[Y]) \text{WSig}^n} \\
\\
\frac{\frac{\Gamma \vdash_l \sigma \text{LSig}^n}{\Gamma \vdash_l \mathbb{L}(\sigma) \quad \Gamma \vdash_l \mathbb{P}(\sigma)} \quad \Delta \vdash_l \sigma \text{LSig}^n \quad \Gamma \vdash \gamma : \Delta}{\Gamma \vdash_l \sigma[Y] \text{LSig}^n \quad \Gamma \vdash \mathbb{L}(\sigma[Y]) \equiv (\mathbb{L}(\sigma))[Y] \quad \Gamma \vdash \mathbb{P}(\sigma[Y]) \equiv (\mathbb{P}(\sigma))[Y]} \\
\\
\frac{}{\Gamma \vdash_l \nu^\bullet \text{LSig}^0} \quad \frac{\Gamma \vdash_l \sigma \text{LSig}^0}{\Gamma \vdash_l \sigma \equiv \nu^\bullet \text{LSig}^0} \quad \frac{\Gamma \vdash_l \sigma \text{LSig}^n \quad \Gamma, A \vdash_l T}{\Gamma, \mathbb{P}(\sigma) \vdash s : A[\mathfrak{p}^1]} \\
\Gamma \vdash_l \nu^+(\sigma, s, T) \text{LSig}^{n+1} \\
\\
\frac{\Theta \vdash_l \sigma \text{LSig}^{n+1} \quad \Gamma \vdash \gamma : \Theta}{\Gamma \vdash_l \nu\pi_1(\sigma) \text{LSig}^n \quad \Gamma \vdash_l \nu\pi'_1(\sigma) \quad \Gamma, \mathbb{P}(\nu\pi_1(\sigma)) \vdash \nu\pi_s(\sigma) : \nu\pi'_1(\sigma)[\pi_1] \quad \Gamma, \nu\pi'_1(\sigma) \vdash_l \nu\pi_2(\sigma)} \\
\Gamma \vdash \nu^+(\nu\pi_1(\sigma), \nu\pi_s(\sigma), \nu\pi_2(\sigma)) \equiv \sigma \text{LSig}^{n+1} \\
\\
\frac{\Gamma \vdash \sigma \text{LSig}^n \quad \Gamma \vdash A \quad \Gamma, \mathbb{P}(\sigma) \vdash s : A[\pi_1] \quad \Gamma, A \vdash_i T}{\Gamma \vdash \nu\pi_1(\nu^+(\sigma, s, T)) \equiv \sigma \text{LSig}^n \quad \Gamma \vdash \nu\pi'_1(\nu^+(\sigma, s, T)) \equiv A} \\
\Gamma, \mathbb{P}(\sigma) \vdash \nu\pi_s(\nu^+(\sigma, s, T)) \equiv s : A[\pi_1] \quad \Gamma, A \vdash \nu\pi_2(\nu^+(\sigma, s, T)) \equiv T \\
\\
\frac{\Gamma \vdash \gamma \text{LSig}^\Theta}{\Gamma \vdash \nu^\bullet[\gamma] \equiv \nu^\bullet \text{LSig}^n \quad \Gamma \vdash (\nu^+(\sigma, s, T))[Y] \equiv \nu^+(\sigma[Y], s[Y^\uparrow], T[Y^\uparrow]) \text{LSig}^{n+1}} \\
\Gamma \vdash (\nu\pi_1(\sigma))[Y] \equiv \nu\pi_1(\sigma[Y]) \text{LSig}^n \quad \Gamma \vdash (\nu\pi'_1(\sigma))[Y] \equiv \nu\pi'_1(\sigma[Y]) \\
\Gamma, \mathbb{P}(\sigma) \vdash \nu\pi_s(\sigma)[Y] \equiv \nu\pi_s(\sigma[Y]) : A[\pi_1] \quad \Gamma, \nu\pi'_1(\sigma[Y]) \vdash (\nu\pi_2(\sigma))[Y^\uparrow] \equiv \nu\pi_2(\sigma[Y]) \\
\\
\frac{}{\Gamma \vdash \mu^\bullet : \mathbb{L}(\nu^\bullet)} \quad \frac{\Gamma \vdash \ell : \mathbb{L}(\nu^\bullet)}{\Gamma \vdash \ell \equiv \mu^\bullet : \mathbb{L}(\nu^\bullet)} \quad \frac{\Gamma \vdash \ell : \mathbb{L}(\sigma)}{\Gamma \vdash \mathbb{P}(\ell) : \mathbb{P}(\sigma)} \quad \frac{\Gamma \vdash \gamma : \Theta}{\Gamma \vdash (\mathbb{P}(\ell))[Y] \equiv \mathbb{P}(\ell[Y]) : \mathbb{P}(\sigma)} \\
\\
\frac{\Gamma \vdash \ell : \mathbb{L}(\sigma) \quad \Gamma, A \vdash t : T}{\Gamma, \mathbb{P}(\sigma) \vdash s : A[\mathfrak{p}^1]} \\
\Gamma \vdash \mu^+(\ell, t) : \mathbb{L}(\nu^+(\sigma, s, T)) \quad \frac{\Gamma \vdash \ell : \mathbb{L}(\sigma)}{\Gamma \vdash \mu^+(\mu\pi_1(\ell), \mu\pi_2(\ell)) \equiv \ell : \mathbb{L}(\sigma)} \\
\\
\frac{\Gamma \vdash \ell : \mathbb{L}(\sigma)}{\Gamma \vdash \mu\pi_1(\ell) : \mathbb{L}(\nu\pi_1(\sigma)) \quad \Gamma, \nu\pi'_1(\sigma) \vdash \mu\pi_2(\ell) : \nu\pi_2(\sigma)} \\
\\
\frac{\Gamma \vdash \ell : \mathbb{L}(\sigma) \quad \Gamma \vdash A \quad \Gamma, \mathbb{P}(\sigma) \vdash s : A[\pi_1] \quad \Gamma, A \vdash t : T}{\Gamma \vdash \mu\pi_1(\mu^+(\ell, t)) \equiv \ell : \mathbb{L}(\sigma) \quad \Gamma, A \vdash \mu\pi_2(\mu^+(\ell, t)) \equiv t : T}
\end{array}$$

$$\begin{array}{c}
\Gamma \vdash \gamma : \Theta \\
\hline
\Gamma \vdash \mu^\bullet[\gamma] \equiv \mu^\bullet : \mathbb{L}(v^\bullet) \quad \Gamma \vdash (\mu^+(\ell, t))[\gamma] \equiv \mu^+(\ell[\gamma], t[\gamma^\uparrow]) : \mathbb{L}(v^+(\sigma, s, T))[\gamma] \\
\Gamma \vdash (\mu\pi_1(\ell))[\gamma] \equiv \mu\pi_1(\ell[\gamma]) : \mathbb{L}(v\pi_1(\sigma))[\gamma] \\
\Gamma, v\pi_1'(\sigma[\gamma]) \vdash (\mu\pi_2(\ell))[\gamma^\uparrow] \equiv \mu\pi_2(\ell[\gamma]) : v\pi_2(\sigma) \\
\\
\Gamma \vdash \mathbb{P}(v^\bullet) \equiv \top \quad \Gamma \vdash \mathbb{P}(v^+(\sigma, s, T)) \equiv \Sigma(\mathbb{P}(\sigma), T[\pi_1, s]) \\
\Gamma \vdash \mathbb{P}(\mu^\bullet) \equiv () : \mathbb{P}(v^\bullet) \quad \Gamma \vdash \mathbb{P}(\mu^+(\ell, t)) \equiv (\mathbb{P}(\ell), t[\pi_1, s][\mathbb{P}^0, \mathbb{P}(\ell)]) : \\
\\
\Gamma \vdash_i A \quad \Gamma, A \vdash_i B \quad \Gamma \vdash_j T \\
\hline
\Gamma \vdash_{i \sqcup j} \text{CaseTy}(A, B, R) \equiv \Pi(A, \Pi(\Pi(B, R[\mathbb{P}^2]), R[\mathbb{P}^2])) \\
\\
\Gamma \vdash_m \tau \text{WSig}^{n+1} \quad \Gamma \vdash_j R \\
\hline
\Gamma \vdash_{m \cup j} \text{RecSig}(\tau, R) \equiv v^+(\text{RecSig}(w^-(\tau), R), \pi_2, \text{CaseTy}(w\pi_1^0(\tau), w\pi_2^0(\tau), R)) \text{LSig}^{n+1}) \\
\\
\Gamma \vdash_m \tau \text{WSig}^0 \quad \Gamma \vdash_j R \quad \Gamma \vdash \tau \text{WSig}^N \quad \Gamma \vdash \ell : \mathbb{L}(\text{RecSig}(\tau, R)) \quad j < N \\
\hline
\Gamma \vdash_{m \cup j} \text{RecSig}(\tau, R) \equiv v^\bullet \text{LSig}^0 \quad \Gamma \vdash R\pi^j(\ell) : (\text{CaseTy}(w\pi_1^j(\tau), w\pi_2^j(\tau), R))[\pi_1] \\
\\
\Gamma \vdash \ell : \mathbb{L}(\text{RecSig}(\tau, T)) \quad \Gamma \vdash t : \text{El}(W(\tau)) \\
\hline
\Gamma \vdash \text{Wrec}(\tau, \ell, t) : T \\
\\
\Gamma \vdash \tau \text{WSig}^N \quad j < N \quad \Gamma \vdash \ell : \mathbb{L}(\text{RecSig}(\tau, R)) \\
\hline
\Gamma \vdash R\pi^j(\ell) : \text{CaseTy}(w\pi_1^j(\tau), w\pi_2^j(\tau), R) \\
\\
\Gamma \vdash R\pi^{n+1}(\ell) \equiv R\pi^n(\mu\pi_1(\ell)) : (\text{CaseTy}(w\pi_1^{n+1}(\tau), w\pi_2^{n+1}(\tau), R))[\pi_1] \\
\Gamma \vdash R\pi^0(\ell) \equiv \mu\pi_2(\ell)[\mathbb{P}^0, \mathbb{P}(\mu\pi_1(\ell))] : (\text{CaseTy}(w\pi_1^0(\tau), w\pi_2^0(\tau), R))[\pi_1] \\
\\
\Gamma \vdash h : \mathbb{L}(\text{RecSig}(\tau, R)) \\
\hline
\Gamma \vdash \text{Wrec}(\tau, h, \text{Wsup}_j(\tau, a, b)) \equiv \text{app}(\text{app}(R\pi^j(h))[(\mathbb{P}^0, a)])[(\mathbb{P}^0, \lambda(\text{Wrec}(\tau, h[\pi_1], b)))] : R \\
\\
\Gamma \vdash h : \sigma_1 \twoheadrightarrow \sigma_2 \quad \Gamma, A_1 \vdash T_1 \quad \Gamma, A_2 \vdash T_2 \quad \Gamma, A_2 \vdash t : T_2 \\
\hline
\Gamma \vdash \text{Identity} : \sigma \twoheadrightarrow \sigma \quad \Gamma \vdash \text{Override}(h, t) : (v^+(\sigma_1, s_1, T_1)) \twoheadrightarrow (v^+(\sigma_2, s_2, T_2)) \\
\\
\Gamma \vdash h : \sigma_1 \twoheadrightarrow \sigma_2 \quad \Gamma, A_2 \vdash t : T \quad \Gamma \vdash h : \sigma_1 \twoheadrightarrow \sigma_2 \quad \Gamma, A_1 \vdash T \quad \Gamma, A_2 \vdash \uparrow^s : A_1[\pi_1] \\
\hline
\Gamma \vdash \text{Extend}(h, t) : \sigma_1 \twoheadrightarrow (v^+(\sigma_2, s_2, T)) \quad \Gamma \vdash \text{Inherit}(h) : v^+(\sigma_1, s_1, T) \twoheadrightarrow v^+(\sigma_2, s_2, T[(\pi_1, \uparrow^s)]) \\
\\
\Gamma \vdash h : \sigma_1 \twoheadrightarrow \sigma_2 \quad \Gamma \vdash \ell : \mathbb{L}(\sigma_1) \quad \Gamma \vdash h : \sigma_1 \twoheadrightarrow \sigma_2 \quad \Gamma, A_2 \vdash \uparrow^s : A_1[\pi_1] \\
\hline
\Gamma \vdash \text{inh}(h, \ell) : \mathbb{L}(\sigma_2) \quad \Gamma, A_2 \vdash h_{\text{inner}} : \tau_1[(\pi_1, \uparrow^s)] \twoheadrightarrow \tau_2 \\
\Gamma \vdash \text{Nest}(h, h_{\text{inner}}) : v^+(\sigma_1, s_1, \mathbb{L}(\tau_1)) \twoheadrightarrow v^+(\sigma_2, s_2, \mathbb{L}(\tau_2)) \\
\\
\Gamma \vdash \text{Identity}[\gamma] \equiv \text{Identity} \quad \Gamma \vdash \text{Override}(h, t)[\gamma] \equiv \text{Override}(h[\gamma], t[\gamma^\uparrow]) \\
\\
\Gamma \vdash \text{Extend}(h, t)[\gamma] \equiv \text{Extend}(h[\gamma], t[\gamma^\uparrow]) \quad \Gamma \vdash \text{Inherit}(h)[\gamma] \equiv \text{Inherit}(h[\gamma]) \\
\\
\Gamma \vdash \text{inh}(h, \ell)[\gamma] \equiv \text{inh}(h[\gamma], \ell[\gamma]) : \mathbb{L}(\sigma_2)[\gamma] \quad \Gamma \vdash \text{Nest}(h, h_{\text{inner}})[\gamma] \equiv \text{Nest}(h[\gamma], h_{\text{inner}}[\gamma^\uparrow])
\end{array}$$

$$\begin{array}{c}
\frac{\Gamma \vdash \ell : \mathbb{L}(\sigma)}{\Gamma \vdash \text{inh}(\text{Identity}, \ell) \equiv \ell : \mathbb{L}(\sigma)} \\
\frac{\Gamma \vdash h : \sigma_1 \rightarrow \sigma_2 \quad \Gamma, \mathbb{P}(\sigma_2) \vdash s_2 : A_2[\pi_1] \quad \Gamma, A_2 \vdash t : T \quad \Gamma \vdash \ell : \mathbb{L}(\sigma_1)}{\Gamma \vdash \text{inh}(\text{Extend}(h, t), \ell) \equiv \mu^+(\text{inh}(h, \ell), t) : \mathbb{L}(v^+(\sigma_1, s_2, T))} \\
\frac{\Gamma \vdash h : \sigma_1 \rightarrow \sigma_2 \quad \Gamma \vdash m : \mathbb{L}(\sigma_1) \quad \Gamma, A_1 \vdash t : T \quad \Gamma, \mathbb{P}(\sigma_2) \vdash s_2 : A_2[\pi_1] \quad \Gamma, A_2 \vdash \uparrow^s : A_1[\pi_1]}{\Gamma \vdash \text{inh}(\text{Inherit}(h), \mu^+(m, t)) \equiv \mu^+(\text{inh}(h, m), t[(\pi_1, \uparrow^s)]) : \mathbb{L}(v^+(\sigma_2, s_2, T[\pi_1, \uparrow^s]))} \\
\frac{\Gamma \vdash h : \sigma_1 \rightarrow \sigma_2 \quad \Gamma \vdash \ell : \mathbb{L}(\sigma_1) \quad \Gamma, A_1 \vdash t_1 : T_1 \quad \Gamma, \mathbb{P}(\sigma_1) \vdash s_1 : A_1[\mathbf{p}^1] \quad \Gamma, A_2 \vdash t_2 : T_2 \quad \Gamma, \mathbb{P}(\sigma_2) \vdash s_2 : A_2[\mathbf{p}^1]}{\Gamma \vdash \text{inh}(\text{Override}(h, t_2), \mu^+(\ell, t_1)) \equiv \mu^+(\text{inh}(h, \ell), t_2) : \mathbb{L}(v^+(\sigma_2, s_2, T_2))} \\
\frac{\Gamma \vdash h : \sigma_1 \rightarrow \sigma_2 \quad \Gamma, A_2 \vdash \uparrow^s : A_1[\pi_1] \quad \Gamma, A_2 \vdash h_{\text{inner}} : \tau_1[(\pi_1, \uparrow^s)] \rightarrow \tau_2}{\Gamma \vdash \text{inh}(\text{Nest}(h, h_{\text{inner}}), \mu^+(\ell, t)) \equiv \mu^+(\text{inh}(h, \ell), \text{inh}(h_{\text{inner}}, t[\pi_1, \uparrow^s])) : \mathbb{L}(v^+(\sigma_1, s_2, \mathbb{L}(\tau_2)))}
\end{array}$$

Linkage transformers as syntactic sugar. As mentioned in Section 5, the five forms of linkage transformers can be thought of as a library of functions that are used to construct linkages from other linkages. In particular, their typing rules are defined in terms of the rest of the typing rules in FMLTT. Thus, linkage transformers $\Gamma \vdash h : \sigma_1 \rightarrow \sigma_2$ can be defined as syntactic sugar via an inductive type (in the metalogic) with four constructors `Extend`(\cdot), `Override`(\cdot), `Inherit`(\cdot), and `Nest`(\cdot). Moreover, `inh`(h, ℓ) and `h`[γ] can be defined as recursive functions (in the metalogic) by induction on this inductive type. We will treat linkage transformers as syntactic sugar in the metatheoretic development in the rest of Section 6.

6.3 A Translation that Compiles Linkages Away

We present a translation from the FMLTT syntax to the fragment of FMLTT without linkage signatures or linkages. The translation echoes how the `FPOP` plugin is implemented (Section 4) as a translation to Coq without facilities for family polymorphism. The translation is type-preserving by construction, as we work in an intrinsically typed setting.

We define the translation $\llbracket \cdot \rrbracket_T$ below; the fragment of the syntax irrelevant to linkages is translated using the identity function and is thus elided. When the context Γ is clear, we use $\llbracket T \rrbracket_T$ to mean $\llbracket \Gamma \vdash T \rrbracket_T$.

We will first define two helpers

- A type $\Gamma \vdash \text{Sig}_j^r \Gamma n$ where $\Gamma \vdash$ and $n \in \mathbb{N}$. The idea is that $\Gamma \vdash \llbracket \Gamma \vdash_j \sigma \text{ LSig}^n \rrbracket_T : \text{Sig}_j^r \Gamma n$
- A function $\mathcal{P} : \{ \sigma' \mid \Gamma \vdash \sigma' : \text{Sig}_j^r \Gamma n \} \rightarrow \{ T \mid \Gamma \vdash T \}$ mapping a term of type $\text{Sig}_j^r \Gamma n$ to a type.

They are defined mutually inductively on n , the signature length.

$$\begin{aligned}
\text{Sig}_j^r \Gamma (n+1) &= \Sigma (\text{Sig}_j^r \Gamma n) (\Sigma \cup_j \\
&\quad (\Sigma (\Pi (\mathcal{P}(\text{var}_0[\mathbf{p}^1]), \text{El}(\text{var}_0[\mathbf{p}^1]))) \\
&\quad \quad \Pi(\text{El}(\text{var}_0[\mathbf{p}^1]), \cup_j))) \\
\text{Sig}_j^r \Gamma 0 &= \top \\
\mathcal{P}(\sigma') &= \Sigma (\text{fst } \sigma') (\text{El}(\text{app}(\text{snd}^3 \sigma')[(\mathbf{p}^1, \text{app}(\text{fst}(\text{snd}^2 \sigma')))])) \\
&\quad \text{when } \Gamma \vdash \sigma' : \text{Sig}_j^r \Gamma n+1 \\
\mathcal{P}(\sigma') &= \top \quad \text{when } \Gamma \vdash \sigma' : \text{Sig}_j^r \Gamma 0
\end{aligned}$$

With the above two helpers, we can define denotation $\llbracket \cdot \rrbracket_T$:

$$\begin{aligned}
\llbracket \Gamma \vdash_j \sigma \text{LSig}^n \rrbracket_T &\in \{ t \mid \Gamma \vdash t : \text{Sig}_j^r \Gamma n \} \\
\Gamma \vdash \mathbb{P}(\sigma) &= \mathcal{P}(\llbracket \sigma \rrbracket_T) \\
\llbracket \Gamma \vdash \mathbb{L}(\sigma) \rrbracket_T &\text{ is defined upon } \llbracket \sigma \rrbracket_T \text{ and inductively on the signature length} \\
\llbracket \Gamma \vdash \mathbb{L}(\sigma) \rrbracket_T &= \llbracket \mathbb{L}(v\pi_1(\sigma)) \rrbracket_T \times \Pi(\text{El}(\text{fst}(\text{snd} \llbracket \sigma \rrbracket_T)))(\text{El}(\text{app}(\text{snd}^3 \llbracket \sigma \rrbracket_T))) \\
&\quad \text{given } \llbracket \Gamma \vdash \sigma \text{LSig}^{n+1} \rrbracket_T \\
\llbracket \Gamma \vdash \mathbb{L}(\sigma) \rrbracket_T &= \top \quad \text{given } \llbracket \Gamma \vdash \sigma \text{LSig}^0 \rrbracket_T \\
\llbracket \Gamma \vdash_j v^\bullet \text{LSig}^0 \rrbracket_T &= () \\
\llbracket \Gamma \vdash_j v^+(\sigma, f, T) \text{LSig}^{n+1} \rrbracket_T &= (\llbracket \sigma \rrbracket_T, c(A), \lambda(f), \lambda(c(T))) \\
\llbracket \Gamma \vdash \mu^\bullet : \mathbb{L}(v^\bullet) \rrbracket_T &= () \\
\llbracket \Gamma \vdash \mu^+ m t : \mathbb{L}(v^+(\sigma, s, T)) \rrbracket_T &= (\llbracket m \rrbracket_T, \lambda(t)) \\
\llbracket \Gamma \vdash \mathbb{P}(m) : \mathbb{P}(\sigma) \rrbracket_T &\text{ is defined upon } \llbracket m \rrbracket_T \text{ and inductively on the signature length} \\
\llbracket \Gamma \vdash \mathbb{P}(m) : \mathbb{P}(\sigma) \rrbracket_T &= () \quad \text{given } \llbracket \Gamma \vdash \sigma \text{LSig}^0 \rrbracket_T \\
\llbracket \Gamma \vdash \mathbb{P}(m) : \mathbb{P}(\sigma) \rrbracket_T &= (\llbracket \mathbb{P}(o) \rrbracket_T, t[(p^1, f)][(\text{id}, \llbracket \mathbb{P}(o) \rrbracket_T)]) \\
&\quad \text{given } \llbracket \Gamma \vdash \sigma \text{LSig}^{n+1} \rrbracket_T, \\
&\quad \text{where } o = v\pi_1(m), t = \text{app}(\text{snd} \llbracket m \rrbracket_T), f = v\pi_s(\sigma)
\end{aligned}$$

The main idea is that a linkage $\Gamma \vdash \mu^+ \text{LSig}^{n+1}$ is translated to a non-dependent tuple while introducing explicit universal quantification to the second component of the tuple; the universal quantification achieves late binding. The translation for \mathbb{P} is given by the relevant β -rules.

We omit validating the equational rules (β , η , and substitution) here. Note that, when we mutually recursively define the type $\text{Sig}^r \Gamma n$ (of signatures $\llbracket \Gamma \vdash t \text{LSig}^n \rrbracket_T$), and the function $\llbracket \mathbb{P} \rrbracket_T = \mathcal{P}$ above, we actually have to prove the two substitution laws $(\mathbb{P}(\sigma))[\gamma] \equiv \mathbb{P}(\sigma[\gamma])$ and $(\text{Sig}^r \Delta n)[\gamma] \equiv \text{Sig}^r \Gamma n$ together.

We have constructed a model for the FMLTT syntax using only the linkage-irrelevant fragment of the syntax. We would have consistency and canonicity for FMLTT immediately, if we could assume the consistency and canonicity of the linkage-irrelevant fragment. However, because our formulation of W-types is unconventional, we choose not to take such an assumption for granted and choose to directly prove consistency and canonicity for FMLTT.

6.4 A Proof Relevant Logical-Relations Model for Canonicity

Now we prove canonicity (and consistency) for FMLTT using a logical-relations model. We follow the reducibility argument of [Kaposi et al. \[2019\]](#), [Coquand \[2019\]](#), and [Sterling \[2019\]](#) to construct our model. The metalanguages of these prior models are based on QITs, categories with families [\[Dybjer 1995\]](#), and the generalized algebraic theory [\[Cartmell 1986\]](#), respectively. Without exposing the reader to too many technical details, our metalanguage should be understood as an instance of any of the above logical frameworks—the difference is that quotienting is manual in our formulation, whereas it is automatic with the logical frameworks.

We state the canonicity theorem first:

THEOREM 6.1 (CANONICITY). *If $\cdot \vdash t : \mathbb{B}$, then either $\cdot \vdash t \equiv \text{tt} : \mathbb{B}$ or $\cdot \vdash t \equiv \text{ff} : \mathbb{B}$.*

Canonicity is a key criterion for a dependent type theory to be considered as a programming language or as a computational foundation for mathematics.

First, we need the mathematical setup to interpret universe levels, following [Sterling \[2019\]](#):

ASSUMPTION 6.1 (SET-THEORETIC UNIVERSE ASSUMPTION). *We assume an infinite hierarchy of Grothendieck universes Set_i for $i \in \mathbb{N}$ in our ambient metalogic.*

We can roughly consider each Grothendieck universe Set_i as the Set_i in Agda:

- Each Set_i is closed under dependent function types and dependent pair types. For example, later, for our interpretation of dependent function types, when we have $\llbracket A \rrbracket_C, \llbracket B \rrbracket_C \in \text{Set}_i$, we will have $\llbracket \Pi(A, B) \rrbracket_C \in \text{Set}_i$.
- The universe hierarchy is cumulative, as $\text{Set}_i \in \text{Set}_{i+1}$ and $\text{Set}_i \subseteq \text{Set}_{i+1}$.
- Thus, if $\llbracket A \rrbracket_C \in \text{Set}_i, \llbracket B \rrbracket_C \in \text{Set}_j$, we will have $\llbracket \Pi(A, B) \rrbracket_C \in \text{Set}_i \cup \text{Set}_j = \text{Set}_{i \cup j}$.

Like most logical-relations proofs, we interpret each judgment and inductively interpret each syntax piece. We are working in an intrinsic setting; thus, even if we omit contexts for brevity, the syntax piece is still well-typed. However, unlike most logical-relations proofs, our logical-relations model is proof-relevant, which is essential for modeling universes properly [Coquand 2019]. Our canonicity model for the base MLTT fragment follows the constructions in Coquand [2019] and Sterling [2019] in that it utilizes the facilitates of the ambient metalogic; for example, we use dependent functions and dependent tuples in the ambient metalogic to model dependent functions and tuples in FMLTT.

$\llbracket \Gamma \vdash k \rrbracket_C$ is a function : $\{ \gamma \mid \cdot \vdash \gamma : \Gamma \} \rightarrow \text{Set}_k$
(i.e., sets indexed by closed substitution)

$\llbracket \Gamma \vdash_j T \rrbracket_C$ is a dependent function : $\prod_{\cdot \vdash \gamma : \Gamma} \prod_{\gamma' \in \llbracket \Gamma \vdash \rrbracket_C(\gamma)}$ $\{ t \mid \cdot \vdash t : T[\gamma] \} \rightarrow \text{Set}_j$

$\llbracket \Gamma \vdash \delta : \Delta \rrbracket_C$ is a dependent function : $\prod_{\cdot \vdash \gamma : \Gamma} \prod_{\gamma' \in \llbracket \Gamma \vdash \rrbracket_C(\gamma)}$ $\llbracket \Delta \vdash \rrbracket_C(\delta \circ \gamma)$

$\llbracket \Gamma \vdash t : T \rrbracket_C$ is a dependent function : $\prod_{\cdot \vdash \gamma : \Gamma} \prod_{\gamma' \in \llbracket \Gamma \vdash \rrbracket_C(\gamma)}$ $\llbracket \Gamma \vdash T \rrbracket_C(\gamma)(\gamma')(t[\gamma])$

$\llbracket \Gamma \vdash T[\sigma] \rrbracket_C(\gamma)(\gamma')(t) = \llbracket T \rrbracket_C(\sigma \circ \gamma)(\llbracket \sigma \rrbracket_C(\gamma)(\gamma'))(t)$

$\llbracket \Gamma \vdash \top \rrbracket_C(\gamma)(\gamma')(t) = \{\star\}$ a singleton set

$\llbracket \Gamma \vdash \perp \rrbracket_C(\gamma)(\gamma')(t) = \emptyset$

$\llbracket \Gamma \vdash \mathbb{B} \rrbracket_C(\gamma)(\gamma')(t) = \begin{cases} \{\star^1\} & \text{if } t \equiv \text{tt} \\ \{\star^2\} & \text{if } t \equiv \text{ff} \\ \emptyset & \text{otherwise} \end{cases}$

$\llbracket \Gamma \vdash \text{Eq}(a, b) \rrbracket_C(\gamma)(\gamma')(t) = \begin{cases} \{\star\} & \text{if } t \equiv \text{refl}(a[\gamma]) \text{ and } a[\gamma] \equiv b[\gamma] \\ \emptyset & \text{otherwise} \end{cases}$

$\llbracket \Gamma \vdash \Pi(A, B) \rrbracket_C(\gamma)(\gamma')(t) = \prod_{\cdot \vdash u : A[\gamma]} \prod_{u' \in \llbracket A \rrbracket_C(\gamma)(\gamma')(u)}$ $\llbracket B \rrbracket_C(\gamma, u)((\gamma', u'))(\text{app}(t)[\text{id}, u])$

$\llbracket \Gamma \vdash \Sigma(A, B) \rrbracket_C(\gamma)(\gamma')(t) = \sum_{u' \in \llbracket A \rrbracket_C(\gamma)(\gamma')(f\text{st } t)}$ $\llbracket B \rrbracket_C(\gamma, \text{fst } t)((\gamma', u'))(\text{snd } t)$

$\llbracket \cdot \vdash \rrbracket_C(\gamma) = \{\star\}$

$\llbracket \Gamma, T \vdash \rrbracket_C(\gamma_t) = \{ (\gamma', t') \mid \gamma' \in \llbracket \Gamma \rrbracket_C(\pi_1 \gamma_t), t' \in \llbracket T[\gamma] \rrbracket_C(\pi_1 \gamma_t)(\gamma')(\pi_2 \gamma_t) \}$

$\llbracket \Gamma \vdash \pi_1 \delta : \Delta \rrbracket_C(\gamma)(\gamma') = \llbracket \delta \rrbracket_C(\gamma)(\gamma')[0]$ get the first element of the tuple

$$\begin{aligned}
& \llbracket \Gamma \vdash \pi_2 \delta : \Delta \rrbracket_C(\gamma)(\gamma') = \llbracket \delta \rrbracket_C(\gamma)(\gamma')[1] \quad \text{get the second element of the tuple} \\
& \llbracket \Gamma \vdash \delta, t : \Delta \rrbracket_C(\gamma)(\gamma') = (\llbracket \delta \rrbracket_C(\gamma)(\gamma'), \llbracket t \rrbracket_C(\gamma)(\gamma')) \\
& \llbracket \Gamma \vdash \text{id} : \Gamma \rrbracket_C(\gamma)(\gamma') = \gamma' \\
& \llbracket \Gamma \vdash \epsilon : \cdot \rrbracket_C(\gamma)(\gamma') = \star \\
& \llbracket \Gamma \vdash \delta_1 \circ \delta_2 : \Delta \rrbracket_C(\gamma)(\gamma') = \llbracket \delta_1 \rrbracket_C(\delta_2 \circ \gamma)(\llbracket \delta_2 \rrbracket_C(\gamma)(\gamma')) \\
& \llbracket \Gamma \vdash t[\sigma] : T[\sigma] \rrbracket_C(\gamma)(\gamma') = \llbracket t \rrbracket_C(\sigma \circ \gamma)(\llbracket \sigma \rrbracket_C(\gamma)(\gamma')) \\
& \llbracket \Gamma \vdash () : \top \rrbracket_C(\gamma)(\gamma') = \star \\
& \llbracket \Gamma \vdash \text{tt} : \mathbb{B} \rrbracket_C(\gamma)(\gamma') = \star^1 \\
& \llbracket \Gamma \vdash \text{ff} : \mathbb{B} \rrbracket_C(\gamma)(\gamma') = \star^2 \\
& \llbracket \Gamma \vdash \text{if}(c, a, b) : T \rrbracket_C(\gamma)(\gamma') = \begin{cases} \llbracket a \rrbracket_C(\gamma)(\gamma') & \text{if } \llbracket c \rrbracket_C(\gamma)(\gamma') = \star^1 \\ \llbracket b \rrbracket_C(\gamma)(\gamma') & \text{if } \llbracket c \rrbracket_C(\gamma)(\gamma') = \star^2 \end{cases} \\
& \llbracket \Gamma \vdash \text{refl}(t) : \text{Eq}(t, t) \rrbracket_C(\gamma)(\gamma') = \star \\
& \llbracket \Gamma \vdash J(w, t) : C[\text{id}, v, t] \rrbracket_C(\gamma)(\gamma') = \llbracket w \rrbracket_C(\gamma)(\gamma') \quad \text{given } \Gamma \vdash t : \text{Eq}(u, v) \\
& \text{Since } \llbracket t \rrbracket_C(\gamma)(\gamma') \text{ witnesses } t[\gamma] \equiv \text{refl}(u[\gamma]) \text{ and } u[\gamma] \equiv v[\gamma] \\
& \llbracket \Gamma \vdash \lambda(t) : \Pi(A, B) \rrbracket_C(\gamma)(\gamma') = \lambda u \lambda u'. \llbracket t \rrbracket_C(\gamma, u)(\gamma', u') \\
& \llbracket \Gamma \vdash \text{app}(t) : B \rrbracket_C(\gamma)(\gamma') = \llbracket t \rrbracket_C(\pi_1 \gamma)(\gamma'[0])(\pi_2 \gamma)(\gamma'[1]) \\
& \llbracket \Gamma \vdash (a, b) : \Sigma(A, B) \rrbracket_C(\gamma)(\gamma') = (\llbracket a \rrbracket_C(\gamma)(\gamma'), \llbracket b \rrbracket_C(\gamma)(\gamma')) \\
& \llbracket \Gamma \vdash \text{fst } t : T \rrbracket_C(\gamma)(\gamma') = \llbracket t \rrbracket_C(\gamma)(\gamma')[0] \quad \text{extract the first element in the tuple} \\
& \llbracket \Gamma \vdash \text{snd } t : T \rrbracket_C(\gamma)(\gamma') = \llbracket t \rrbracket_C(\gamma)(\gamma')[1] \\
& \llbracket \Gamma \vdash_{j+1} \cup_j \rrbracket_C(\gamma)(\gamma')(T) = \{ t \mid \cdot \vdash t : \text{El}(T) \} \rightarrow \text{Set}_j \\
& \llbracket \Gamma \vdash c(T) : \cup_j \rrbracket_C(\gamma)(\gamma') = \llbracket T \rrbracket_C(\gamma)(\gamma') \\
& \llbracket \Gamma \vdash_j \text{El}(T) \rrbracket_C(\gamma)(\gamma')(t) = \llbracket T \rrbracket_C(\gamma)(\gamma')(t)
\end{aligned}$$

Here, \star, \star^1, \star^2 are just some arbitrary fixed elements.

Given the above model for MLTT, there should be a function Π^c such that $\Pi^c(\llbracket A \rrbracket_C, \llbracket B \rrbracket_C) = \llbracket \Pi(A, B) \rrbracket_C$. Type-theoretically speaking, this Π^c uses the *internal dependent function type* of the above model. We hope to use this function when defining the logical-relations model for the rest of FMLTT. However, such a function Π^c is not yet possible because the definition $\llbracket \Pi(A, B) \rrbracket_C$ is not based solely on $\llbracket A \rrbracket_C$ and $\llbracket B \rrbracket_C$, but also on the syntax $\Gamma \vdash A$ and $\Gamma, A \vdash B$.

Thus, we define a new denotation $\llbracket S \rrbracket_C^\bullet := (S, \llbracket S \rrbracket_C)$ that also returns the syntax piece S .⁵ Then, we can have a function Π^\bullet such that $\Pi^\bullet(\llbracket A \rrbracket_C^\bullet, \llbracket B \rrbracket_C^\bullet) = \llbracket \Pi(A, B) \rrbracket_C^\bullet$ now that the syntax is available. Similarly, there are functions $\Sigma^\bullet, (a, \bullet b)$, and $\gamma, \bullet t$ for dependent pair types, dependent pairs, and substitution extension (and more for other constructions). Furthermore, given S^\bullet (i.e., the syntax and its semantic interpretation), we use $(S^\bullet)^c$ to mean the latter of the two.

We need more *internal type-theoretic* constructions:

$$\begin{aligned}
& \text{We define } \text{Con}_k^\bullet := \sum_{\Gamma \vdash k} \{ \gamma : \cdot \vdash \gamma : \Gamma \} \rightarrow \text{Set}_k \\
& \text{and } \text{Ty}_j^\bullet \Gamma^\bullet := \sum_{\Gamma \vdash j T} \prod_{\cdot \vdash \gamma : \Gamma} \prod_{\gamma' \in (\Gamma^\bullet)^c(\gamma)} \{ t : \cdot \vdash t : T[\gamma] \} \rightarrow \text{Set}_j \text{ for } \Gamma^\bullet \in \text{Con}_k^\bullet \\
& \text{and } \text{Tm}^\bullet \Gamma^\bullet T^\bullet := \sum_{\Gamma \vdash t : T} \prod_{\cdot \vdash \gamma : \Gamma} \prod_{\gamma' \in (\Gamma^\bullet)^c} (T^\bullet)^c(\gamma)(\gamma')(t[\gamma]) \\
& \text{and } \text{Sub}^\bullet \Gamma^\bullet \Delta^\bullet := \sum_{\Gamma \vdash \delta : \Delta} \prod_{\cdot \vdash \gamma : \Gamma} \prod_{\gamma' \in (\Gamma^\bullet)^c(\gamma)} (\Delta^\bullet)^c(\delta \circ \gamma)
\end{aligned}$$

⁵This is also called glued interpretation in [Sterling \[2019\]](#).

These four sets are collecting the glued interpretation. For each well-formed type $\Gamma \vdash_j T$, we have its denotation $\llbracket T \rrbracket_C^\bullet \in \text{Ty}_j^\bullet \llbracket \Gamma \rrbracket_C^\bullet$; for each well-typed term $\Gamma \vdash t : T$, we have its denotation $\llbracket t \rrbracket_C^\bullet \in \text{Tm}^\bullet \llbracket \Gamma \rrbracket_C^\bullet \llbracket T \rrbracket_C^\bullet$; etc. These structures should remind the reader of the ones appearing at the beginning of the MLTT canonicity model.

Notice that $\prod_{\vdash \gamma : \Gamma} \prod_{\gamma' \in (\Gamma^\bullet)^\epsilon(\gamma)}$ is part of Ty^\bullet , Tm^\bullet , and Sub^\bullet . A useful fact about them is : given a pair of arbitrary $\cdot \vdash \gamma : \Gamma$ and $\gamma' \in (\Gamma^\bullet)^\epsilon(\gamma)$, we can consider (γ, γ') as an element of $\text{Sub}^\bullet \cdot \Gamma^\bullet$, and vice versa. Thus, we consider the pair (γ, γ') the equivalent form of an element $\gamma^\bullet \in \text{Sub}^\bullet \cdot \Gamma^\bullet$.

Then we define $\text{WSig}_j^{C_n} \Gamma^\bullet := \text{Vector}^n \sum_{A^\bullet \in \text{Ty}_j^\bullet \Gamma^\bullet} \text{Ty}_j^\bullet(\Gamma^\bullet, A^\bullet)$ for $\Gamma^\bullet \in \text{Con}_k^\bullet$, a length- n list of pairs of types. As before, we can define glued interpretation $\text{WSig}_j^{C_n} \Gamma^\bullet := \{ \tau \mid \cdot \vdash \tau \text{ WSig}^n \} \times \text{WSig}_j^{C_n} \Gamma^\bullet$. This will be useful when we interpret judgments for W -type signatures.

Now we can extend the model above to include FMLTT constructs. The key idea of this extension is similar to that of the syntactic translation: we interpret linkage types using the canonicity model of Σ -types (developed in the prior work). We use the inductive facility of the ambient metalogic to justify our W -types.

$\llbracket \Gamma \vdash \sigma \text{LSig}^n \rrbracket_C$ is a list of 3-tuple of length n

$$\llbracket \Gamma \vdash_j \tau \text{WSig}^n \rrbracket_C : \text{WSig}_j^{C_n} \llbracket \Gamma \rrbracket_C^\bullet$$

i.e., a list of 2-tuple of length n

we interpret $\llbracket () \rrbracket, \llbracket \mathbb{P}() \rrbracket$ by induction on the input linkage signature, via L and P^\bullet

$$L(\text{nil}) = \llbracket \top \rrbracket_C$$

$$L((A, s, T) :: tl)(\gamma)(\gamma')(t) = L(tl)(\gamma)(\gamma')(\mu\pi_1(t)) \times (\Pi^\bullet(A, T))^\epsilon(\gamma)(\gamma')(\lambda(\mu\pi_2(t)))$$

$$P^\bullet(\text{nil}) = (\mathbb{P}(v^\bullet), \llbracket \top \rrbracket_C)$$

$$P^\bullet((A, s, T) :: tl) = \Sigma^\bullet(P^\bullet(tl), T[(p^1)^\bullet, s]^\bullet) \quad (\text{doing substitution on } T)$$

$$\llbracket \Gamma \vdash L(\sigma) \rrbracket_C = L(\llbracket \sigma \rrbracket_C)$$

$$\llbracket \Gamma \vdash \mathbb{P}(\sigma) \rrbracket_C = (P^\bullet(\llbracket \sigma \rrbracket_C))^\epsilon \quad \text{discard syntax info}$$

$$\llbracket \Gamma \vdash v^\bullet \text{LSig}^0 \rrbracket_C = \text{nil}$$

$$\llbracket \Gamma \vdash v^+(\sigma, s, T) \text{LSig}^{n+1} \rrbracket_C = (\llbracket A \rrbracket_C^\bullet, \llbracket s \rrbracket_C^\bullet, \llbracket T \rrbracket_C^\bullet) :: \llbracket \sigma \rrbracket_C \quad \text{given } \Gamma, \mathbb{P}(\sigma) \vdash s : A$$

$\llbracket \Gamma \vdash \sigma[\gamma] \text{LSig}^n \rrbracket_C$ is done by point-wise/component-wise substitution

$$\llbracket \Gamma \vdash v\pi_1(\sigma) \text{LSig}^n \rrbracket_C = \text{tl } \llbracket \sigma \rrbracket_C$$

$$\llbracket \Gamma \vdash v\pi'_1(\sigma) \rrbracket_C = ((\text{hd } \llbracket \sigma \rrbracket_C)[0])^\epsilon \quad \text{take the first element in the tuple, ...}$$

$$\llbracket \Gamma, \mathbb{P}(v\pi_1(\sigma)) \vdash v\pi_s(\sigma) : v\pi'_s(\sigma) \rrbracket_C = ((\text{hd } \llbracket \sigma \rrbracket_C)[1])^\epsilon \quad \text{... and discard syntax}$$

$$\llbracket \Gamma, v\pi'_1(\sigma) \vdash v\pi_2(\sigma) \rrbracket_C = ((\text{hd } \llbracket \sigma \rrbracket_C)[2])^\epsilon$$

$$\llbracket \Gamma \vdash \mu^\bullet : \mathbb{L}(v^\bullet) \rrbracket_C = \llbracket () \rrbracket_C$$

$$\llbracket \Gamma \vdash \mu^+(\ell, t) : \mathbb{L}(v^+(\sigma, s, T)) \rrbracket_C(\gamma)(\gamma') = (\llbracket \ell \rrbracket_C(\gamma)(\gamma'), (\lambda^\bullet(\llbracket t \rrbracket_C^\bullet))^\epsilon(\gamma)(\gamma'))$$

$$\llbracket \Gamma \vdash \mathbb{P}(\ell) : \mathbb{P}(\sigma) \rrbracket_C = \llbracket () \rrbracket_C \quad \text{when } \Gamma \vdash \sigma \text{LSig}^0$$

$$\llbracket \Gamma \vdash \mathbb{P}(\ell) : \mathbb{P}(\sigma) \rrbracket_C = ((\ell'^\bullet, \llbracket \mu\pi_2(\ell) \rrbracket_C^\bullet [(p^1)^\bullet, s]^\bullet [(id)^\bullet, \ell'^\bullet]^\bullet))^\epsilon$$

where $\ell'^\bullet = \llbracket \mathbb{P}(\mu\pi_1(\ell)) \rrbracket_C^\bullet$

when $\llbracket \Gamma \vdash \sigma \text{LSig}^{n+1} \rrbracket_C = (A, s, T) :: _$

In this definition, $(_, \bullet _)$ denotes the glued interpretation of dependent tuple construction, $[_ \bullet _]$ denotes the glued interpretation of substitution extension, and $_[_]^\bullet$ denotes the glued interpretation of applying a substitution.

The colored boxes indicate the scopes of the two instances of $_[_]^\bullet$. We use colored boxes here and below to help the reader parse nested glued interpretations.

$$\begin{aligned} \llbracket \Gamma \vdash \mu\pi_1(\ell) : \mathbb{L}(\sigma) \rrbracket_C(\gamma)(\gamma') &= \llbracket \ell \rrbracket_C(\gamma)(\gamma')[0] \quad \text{take the first element in the tuple} \\ \llbracket \Gamma, \nu\pi'_1(\sigma) \vdash \mu\pi_2(\ell) : T \rrbracket_C(\gamma_+)(\gamma'_+) &= ((\llbracket \ell \rrbracket_C(\pi_1\gamma_+)(\gamma'_+[0]))[1])(\pi_2\gamma_+)(\gamma'_+[1]) \end{aligned}$$

$$\llbracket \Gamma \vdash w^\bullet \text{WSig}^0 \rrbracket_C = \text{nil}$$

$$\llbracket \Gamma \vdash w^+(\tau, A, B) \text{WSig}^{n+1} \rrbracket_C = (\llbracket A \rrbracket_C^\bullet, \llbracket B \rrbracket_C^\bullet) :: \llbracket \tau \rrbracket_C$$

$$\llbracket \Gamma \vdash \tau[\gamma] \text{WSig}^n \rrbracket_C \text{ is done by point-wise/component-wise substitution}$$

$$\llbracket \Gamma \vdash w\pi_1^j(\tau) \rrbracket_C = ((j\text{-th element of } \llbracket \tau \rrbracket_C)[0])^c$$

$$\llbracket \Gamma, w\pi_1^j(\tau) \vdash w\pi_2^j(\tau) \rrbracket_C = ((j\text{-th element of } \llbracket \tau \rrbracket_C)[1])^c$$

$$\llbracket \Gamma \vdash w^-(\tau) \text{WSig}^n \rrbracket_C = \text{tl } \llbracket \tau \rrbracket_C$$

$$\llbracket \Gamma \vdash W(\tau) : \mathbb{U} \rrbracket_C(\gamma)(\gamma')(t) = W^C(\llbracket \tau \rrbracket_C[\gamma^\bullet]^\bullet) t$$

$$\llbracket \Gamma \vdash \text{Wsup}_i(\tau, a, b) : \text{El}(W(\tau)) \rrbracket_C(\gamma)(\gamma') = W^C \text{sup } i \quad (\llbracket a \rrbracket_C^\bullet[\gamma^\bullet]^\bullet) \quad (\llbracket b \rrbracket_C^\bullet[\gamma^\uparrow]^\bullet)$$

$$\llbracket \Gamma \vdash \text{CaseTy}(A, B, R) \rrbracket_C = (\Pi^\bullet(\llbracket A \rrbracket_C^\bullet, \Pi^\bullet(\Pi^\bullet(\llbracket B \rrbracket_C^\bullet, \llbracket R \rrbracket_C^\bullet[(p^2)^\bullet]^\bullet), \llbracket R \rrbracket_C^\bullet[(p^2)^\bullet]^\bullet)))^c$$

The interpretation of $\text{Wsup}_i(\tau, a, b)$ is defined via $W^C \text{sup}$, which is in turn defined below. This construction takes three arguments as input: i , a glued $(\llbracket a \rrbracket_C^\bullet[\gamma^\bullet]^\bullet)$ and a glued $(\llbracket b \rrbracket_C^\bullet[\gamma^\uparrow]^\bullet)$. As mentioned earlier, γ^\bullet is an equivalent form of (γ, γ') .

we define $\text{RecSig}(\cdot)$ by induction on the signature, via RS

$$RS \text{ nil } R = \llbracket v^\bullet \rrbracket_C$$

$$RS ((A, B) :: tl) R = v^{+\bullet}(RS \text{ tl } R, \pi_2^\bullet, \text{CaseTy}^\bullet(A, B, R))$$

$$\llbracket \Gamma \vdash \text{RecSig}(\tau, R) \rrbracket_C = (RS \llbracket \tau \rrbracket_C \llbracket R \rrbracket_C)^c$$

$$\llbracket R\pi^j(\ell) \rrbracket_C = \text{take the } j\text{-th field from } \ell$$

$$\begin{aligned} \llbracket \Gamma \vdash W\text{rec}(\tau, \ell, t) : T \rrbracket_C(\gamma)(\gamma') &= W^C \text{rec} \quad \llbracket \tau \rrbracket_C^\bullet[\gamma^\bullet]^\bullet \\ &\quad (\lambda w. \llbracket R \rrbracket_C(\gamma)(\gamma')(W\text{rec}(\tau[\gamma], \ell[\gamma], w))) \\ &\quad f^r \\ &\quad t[\gamma] \\ &\quad (\llbracket t \rrbracket_C(\gamma)(\gamma')) \end{aligned}$$

$$\text{where } f^r \text{ j } a^\bullet b^\bullet = \text{let } \rho^\bullet \in \text{Tm}^\bullet \quad (\cdot^\bullet, \cdot^\bullet \text{ B}^\bullet[\text{id}^\bullet, \cdot^\bullet a^\bullet]^\bullet) \quad (R^\bullet[\gamma^\bullet]^\bullet[(p^1)^\bullet]^\bullet)$$

$$\text{s.t. } \rho^\bullet := (W\text{rec}(\tau, \ell[\gamma \circ p^1], b), b^c) \text{ in}$$

$$(\text{app}^\bullet(\text{app}^\bullet(\text{R}\pi^{j^\bullet}(\llbracket \ell \rrbracket_C^\bullet[\gamma^\bullet]^\bullet)) [\text{id}^\bullet, \bullet a^\bullet]^\bullet) [\text{id}^\bullet, \bullet \lambda^\bullet(\rho^\bullet)]^\bullet)^c (\epsilon)(\star)$$

given $\Gamma \vdash \ell : \text{RecSig}(\tau, \text{R})$

where $R^\bullet = \llbracket \text{R} \rrbracket_C^\bullet, B^\bullet = \text{w}\pi_2^{j^\bullet}(\llbracket \tau \rrbracket_C^\bullet)$

The semantic interpretations above are defined together with the following inductively defined indexed set W^C (with only one constructor $W^C \text{sup}$)

(Note: we also use $(a \in A \rightarrow B(a))$ as another notation for dependent function $\prod_{a \in A} B(a)$.)

$$\begin{aligned} \text{Inductive } W^C : (\tau^\bullet \in \text{WSig}_i^N \llbracket \cdot \rrbracket_C^\bullet) &\rightarrow \{t \mid \cdot \vdash t : \text{El}(W(\tau))\} \rightarrow \text{Set}_{i+1} \text{ where} \\ W^C \text{sup} : j < N \rightarrow a^\bullet \in \text{Tm}^\bullet \cdot \text{w}\pi_1^{j^\bullet}(\tau^\bullet) & \\ &\rightarrow b^\bullet \in (\text{Tm}^\bullet (\cdot, \bullet \text{w}\pi_2^{j^\bullet}(\tau^\bullet) [\text{id}^\bullet, a^\bullet]^\bullet) \text{El}^\bullet(W^\bullet(\tau^\bullet)) [(p^1)^\bullet]^\bullet) \\ &\rightarrow W^C \tau^\bullet \text{Wsup}_j(\tau, a, b) \end{aligned}$$

and its eliminator $W^C \text{rec}$

$$\begin{aligned} W^C \text{rec} : (\tau^\bullet \in \text{WSig}_i^N \llbracket \cdot \rrbracket_C^\bullet) &\rightarrow (P : \{t \mid \cdot \vdash t : \text{El}(W(\tau))\} \rightarrow \text{Set}_k) \\ &\rightarrow \left(j < N \rightarrow a^\bullet \in \text{Tm}^\bullet \cdot \text{w}\pi_1^{j^\bullet}(\tau^\bullet) \right. \\ &\quad \rightarrow \left\{ b \mid (\cdot, \text{w}\pi_2^{j^\bullet}(\tau) [\text{id}, a]) \vdash b : \text{El}(W(\tau)) [p^1] \right\} \\ &\quad \rightarrow \left(\gamma^\bullet \in \text{Sub}^\bullet \cdot (\cdot, \bullet \text{w}\pi_2^{j^\bullet}(\tau^\bullet) [\text{id}^\bullet, \bullet a^\bullet]^\bullet) \rightarrow P(b[\gamma]) \right) \\ &\quad \left. \rightarrow P(\text{Wsup}_j(\tau, a, b)) \right) \\ &\rightarrow \cdot \vdash t : \text{El}(W(\tau)) \rightarrow W^C \tau^\bullet t \rightarrow P t \\ W^C \text{rec } \tau^\bullet P f t (W^C \text{sup } a^\bullet b^\bullet) &= f a^\bullet b (\lambda \gamma^\bullet. W^C \text{rec } \tau^\bullet P f (b[\gamma]) ((b^\bullet)^c \gamma^\bullet)) \end{aligned}$$

Note that in $W^C \text{sup}$, the b^\bullet uses the definition of $\llbracket W(\tau) \rrbracket_C$, which after unfolding, recursively references W^C in a strictly positive position. We do not distinguish (b, b^c) and b^\bullet for simplicity.

The idea of the proof of W -type is, as mentioned, mirroring the facility of the inductive type in the ambient logic into FMLTT. The main difference between W^C in the ambient logic and $W(\cdot)$ in the FMLTT, is that W^C is only witnessing those reducible *closed terms*. Thus when using W^C to model $W(\cdot)$, we need to do closed substitution properly.

Again, we omit validating the equational rules (β , η , and substitution) here.

We state the fundamental property of the logical-relations model.

THEOREM 6.2 (FUNDAMENTAL PROPERTY). *If $\Gamma \vdash t : T$, then its semantic interpretation is a dependent function such that $\llbracket t \rrbracket_C : \prod_{\vdash \gamma : \Gamma} \prod_{\gamma' \in \llbracket \Gamma \rrbracket_C(\gamma)} \llbracket \Gamma \vdash T \rrbracket_C(\gamma)(\gamma')(t[\gamma])$.*

The first consequence of this model is the consistency of FMLTT—we cannot derive $\cdot \vdash t : \perp$. Otherwise, we would have an element in the empty set, $\llbracket \cdot \vdash t : \perp \rrbracket_C(\epsilon)(\star) \in \llbracket \perp \rrbracket_C(\epsilon)(\star)(t[\gamma]) = \emptyset$, a contradiction.

THEOREM 6.3 (CONSISTENCY). *The typing judgment $\cdot \vdash t : \perp$ is not derivable for any term t .*

Next, with the logical-relations model, we can map an arbitrary closed boolean term $\cdot \vdash t : \mathbb{B}$ to get the result $\llbracket \cdot \vdash t : \mathbb{B} \rrbracket_C(\epsilon)(\star) = \star^1$ or \star^2 , witnessing the proof of $t \equiv \text{tt}$ or $t \equiv \text{ff}$ by the definition of our model, arriving at [Theorem 6.1](#).

Further, with the help of eta rules, we have the following canonical forms.

THEOREM 6.4 (CANONICAL FORMS).

- If $\cdot \vdash t : \text{El}(W(\tau))$ and $\cdot \vdash \tau \text{ WSig}^n$, then $\cdot \vdash t \equiv \text{Wsup}_j(\tau, a, b) : \text{El}(W(\tau))$ for some $\cdot \vdash a : A$, $B[(\text{id}, a)] \vdash b : \text{El}(W(\tau))$, and $j < n$
- If $\cdot \vdash t : \mathbb{B}$ then $\cdot \vdash t \equiv \text{tt} : \mathbb{B}$ or $\cdot \vdash t \equiv \text{ff} : \mathbb{B}$
- If $\cdot \vdash t : \mathbb{L}(\sigma)$ with $\cdot \vdash \sigma \text{ LSig}^n$, then $\cdot \vdash t \equiv \mu^+(o, t) : \mathbb{L}(\sigma)$ for some $\cdot \vdash o : \mathbb{L}(\mu\pi_1(\sigma))$ and $\nu\pi_2(\sigma) \vdash t : \mu\pi_2(\sigma)$
- If $\cdot \vdash t : \Sigma(A, B)$ then $\cdot \vdash t \equiv (a, b) : \Sigma(A, B)$ with $\cdot \vdash a : A$ and $\cdot \vdash b : B[(\text{id}, a)]$ ⁶

6.5 Using FMLTT’s Linkage Transformers to Model a Derived Family

Here, we sketch how to use the “library” of linkage transformers in FMLTT to inductively construct a linkage transformer that models a derived family of STLC.

We use `STLCBool` as an example. `STLCBool` extends `STLC` with boolean values (`tm_true` and `tm_false`) and if-then-else expressions. The left column of the table below shows initial code excerpted from this family. Each cell in the last column defines a linkage transformer h_i inductively constructed from the linkage transformers h_0, \dots, h_{i-1} using one of the introduction forms. The goal is to eventually construct a linkage transformer h_n representing the entire family `STLCBool`.

Most steps are self-explanatory. Of note are the two grayed rows. They are constructing a linkage transformer h_β containing the case handlers for `subst`. This h_β is then appended to h_3 as a nested linkage transformer.

surface-syntax program	i	$\cdot \vdash h_i : \sigma_i \rightarrow \sigma'_i$
<code>Family STLCBool extends STLC.</code>	0	Identity
<code>FInductive tm +=</code>	1	Override($h_0, W(\tau'_{\text{tm}})$)
<i>(* existing constructors *)</i>	2	Override($h_1, \text{Wsup}(\tau'_{\text{tm}}, \top, \perp)$)
<code> tm_true ...</code>	3	Extend($h_2, \text{Wsup}(\tau'_{\text{tm}}, \top, \perp)$)
<code>FRecursion subst ... +=</code>	α	Identity
<code>Case tm_true := ... Case ...</code>	β	Extend(h_α, \dots)
<code>End subst.</code>	4	Nest(h_3, h_β)
	5	Override($h_4, \lambda t. \text{Wrec}(\tau'_{\text{tm}}, t, \text{substCases})$)
<code>FInductive ty += ...</code>	6	...
<i>(* Inherit env *)</i>	7	Inherit(h_6)
<i>(* Inherit empty *)</i>	8	Inherit(h_7)
...

⁶We emphasize the last one because $\mathbb{P}(\ell)$ is a dependent pair


7 CASE STUDIES

Type safety of STLCs. The first case study is the mechanization of the type safety theorem of STLC and those of its extensions, which has been occurring in the examples in this paper. The code base is ported from Software Foundations [Pierce et al. 2022]. The linguistic nature of our approach allows us to retain a programming style similar to the original proofs in Software Foundations.

The base STLC family consists of ~360 LOC, about the same as an STLC development not using FPOP. Lines of code in each of the four derived families (Y , \times , $+$, and μ in the Venn diagram) vary from 100 to 250, largely depending on how many constructors they add to the inductive types. Without FPOP, the same STLC code would have to be duplicated for each feature.

Using individual families to organize the mechanization of individual language features leads to a modular design that also facilitates code reuse. Individually developed features can be easily composed (as mixins) to form new STLC variants (e.g., $Y+\mu$). Such a feature composition often requires only a few lines of code.

Composing features can lead to *feature interactions* [Batory et al. 2011]: features working correctly in isolation may require coordination when composed. For example, composing \times and μ (Figure 3) creates an obligation to extend `tysubst` to handle `ty_prod`, which the type-checker enforces. Composing contradictory features (e.g., a fixpoints construct and a strong normalization theorem) would lead to unprovable proof obligations.




Elimination of inductive types defined via `FInductive` is mostly via the `FRecursion` and `FInduction` commands. An exception is a handful of trivial “inversion lemmas”. For example, consider the lemma $\forall t, \neg \text{step } \text{tm_true } t$ stating that `tm_true` is irreducible. If `step` were an ordinary inductive type, then it could be proved in Coq simply by `intros t H; inversion H`. But `step` is extensible. So one way to prove the lemma is by `FInduction` on `step` and verifying that a derived family does not accidentally make `tm_true` reducible. We observe that it is lighter-weight to use overriding (Section 3.3) instead: the programmer can specify that the proof of the lemma should be overridden in any derived family that further binds `step`, and in return, they are permitted to treat `step` as an ordinary inductive type in the proof and thus use `inversion` to prove it. The plugin then automatically tries the same proof script in a derived family to override the proof. Although proof scripts rather than proof terms are reused, this practice seems justified by the triviality of the lemmas and the terseness of the proof scripts.

Abstract interpreters for imperative languages. Our second case study is a mechanization of abstract interpreters for simple imperative languages. In addition to a soundness proof, this case study produces abstract interpreters that are directly ready for program extraction.

The code is organized into four families. A base family `Imp` (~200 LOC) defines via `FInductive` the abstract syntax of a while-language with pure expressions and impure statements. The semantics is given by an interpreter defined as a CEK-style abstract machine [Felleisen and Friedman 1986] and parameterized by a fuel value. Family `Imp` defines the interpreter via `FRecursion`.

A second family `ImpGAI` (~550 LOC) extends `Imp`. It exports a generic framework for deriving abstract interpreters with partial-correctness guarantees. Soundness of the abstract interpreter, `analyze`, is stated with respect to the interpreter, `eval`, inherited from `Imp`. The theorem says that the concretization relation `RState` over a concrete state `S` and an abstract state `absS` is preserved by the analysis:



$\forall \text{ stmt fuel } S \text{ absS}, R\text{State } S \text{ absS} \rightarrow R\text{State} (\text{eval fuel stmt } S) (\text{analyze fuel stmt absS})$

analyze is defined via **FRecursion**, and the soundness theorem is proved via **FInduction**. This family leaves fields representing the abstract domain, the concretization relation, monotonicity of transfer functions, etc. largely unspecified or unproven—a derived family can further bind these “parameters” by overriding appropriate fields (and also possibly extend the abstract syntax), to create a sound, runnable abstract interpreter for a (possibly extended) while-language.

The next two families both extend ImpGAI. Family ImpTI (~200 LOC) is an abstract interpreter doing type inference [Cousot 1997]. Family ImpCP (~300 LOC) extends the abstract syntax with natural-number arithmetic, and further binds the generic abstract interpreter to perform constant propagation.

Our implementation of family polymorphism is compatible with Coq’s program extraction feature. We extract the two verified abstract interpreters to OCaml. Testing the extracted program over simple queries returns expected results.

In addition to the two case studies above, we also use extensible inductive types for modeling extensible context-free grammars and derive decision procedures for language membership.

8 RELATED WORK

Approaches to modular mechanization or proof reuse exist, with different focuses and trade-offs.

Encodings based on product lines or DTC. Delaware et al. [2011] engineer product lines of theorems and proofs built from feature modules. Feature composition is manual, which seems to have motivated later approaches based on data types à la carte (DTC) [Swierstra 2008].

The original DTC encoding requires type-level general recursion that fails the strict positivity check imposed by proof assistants including Coq and Agda. Delaware et al. [2013a] introduce meta-theory à la carte (MTC): it overcomes the problem by using Church encodings for data types and using Mendler-style folds for evaluation, though it requires Set impredicativity. Feature composition is automated through heavy use of type classes. The framework is implemented as a Coq library. Schwaab and Siek [2013] adapt DTC to Agda by considering a restricted class of functors that admit least fixed points. Keuchel and Schrijvers [2013] use datatype-generic programming techniques for the underlying representation of type-level fixed points and avoid Set impredicativity.

All of these approaches are largely extralinguistic, in that they work within the confine of the language offered by a proof assistant, which comes with trade-offs. On the one hand, they can be conveniently distributed as libraries, and the encoding can be more easily adapted for new purposes. For example, MTC has been applied to implementing composable program adverbs [Li and Weirich 2022] and been adapted to allow feature extensions, such as reference cells and exceptions, that require type changes [Delaware et al. 2013b; van der Rest et al. 2022].

On the other hand, the extralinguistic nature of the approaches tends to lead to non-idiomatic code and offset their user-friendliness. In particular, because data types have to be encoded (rather than expressed through natively supported inductive types), the resulting code can be obtuse at first blush, making the programming style inaccessible to non-experts. In addition, extra programmer effort may be required, such as having to manually prove additional well-formedness conditions.

Forster and Stark [2020] introduce *Coq à la carte*. It still follows DTC, but rather than embracing DTC's use of generic fixed points, it considers specific instantiations instead. The resulting mechanism appears more streamlined than prior à la carte approaches particularly for its extensive tool support for generating boilerplate code. But even with the tool support, components (e.g., `subst1em`) of individually developed feature extensions have to be composed separately by invoking the tool.

Our approach addresses the expression problem by extending the linguistic facilities offered by a proof assistant. Families, in particular, offer an organizational advantage. They allow grouping and coevolving related types, functions, and proofs without explicit parameterization; all further-bindable fields are automatically extensibility hooks. Because family polymorphism does not require explicit parameterization or complex encodings, the resulting programming experience and code are accessible to the working Coq programmer. The more OO aspects of family polymorphism, such as the ability to use families as mixins and the ability to grow a series of mechanized languages in integral increments, also facilitate extensibility and reuse with minimal programmer effort.

Proof reuse and proof repair. Boite [2004] addresses proof reuse specifically in response to inductive types extended with new constructors. Proof reuse is via a tactic that adapts the original proof to the extended inductive type while generating proof obligations, so rechecking of proof terms is entailed. The design requires distinct names for a base inductive type and its extensions (including distinct names for constructors), while FPOP allows names to be late bound.

Mulhern [2006] introduces a heuristic approach that allows proofs for multiple small languages to be combined to yield proofs for composite languages, as long as the proof structure follows the same pattern. Johnsen and Lüth [2004] enable proof reuse in Isabelle by adapting theorems

from one setting for reuse in another: proof terms are transformed by first explicitly stating all assumptions, and then abstracting over function symbols and type constants.

Pumpkin Pi [Ringer et al. 2021] is a Coq plugin that helps repair proofs broken by changes in type definitions. Its decompiler from proof terms to proof scripts prioritizes suggesting useful tactics over soundness. While Pumpkin Pi focuses on refactoring existing proofs in response to changed definitions, our solution can be viewed as an effort to preempt refactoring by enabling the programmer to write code that has built-in hooks for future extensions.

The expression problem. Solutions abound. Almost all involve some form of either explicit or implicit parameterization as extensibility hooks. Our approach is the first that applies family polymorphism [Ernst 2001], mostly seen in OO languages, to the context of mechanized proofs.

Blume et al. [2006] address the expression problem for a core subset of Standard ML by combining explicitly coded open recursion with a design that allows pattern-matching cases to be defined separately and combined later. Our `FRecursion` and `FInduction` commands achieve a similar functionality, with families making open recursion implicit and bestowing organizational power.

ML-style modules, like families, are a modularity mechanism, but with a focus on abstraction rather than extensibility. Both FMLTT and the module system of Stone and Harper [2000] use singletons to model and control the propagation of definitions. MixML [Rossberg and Dreyer 2013] integrates mixins into ML and handles the idiosyncrasies of ML modules, while our work supports mixins in the presence of extensible inductive types.

9 CONCLUSION

It is hard to write modular, extensible code and proofs. We have presented a solution that equips a proof assistant with linguistic facilities for family polymorphism. The language design ensures that the expressive power brought by family polymorphism is in harmony with the strictness of a proof assistant, while incurring low cognitive overhead and allowing an idiomatic programming experience. We implement the design via a translation to Coq and demonstrate its applicability using case studies. A novel dependent type theory formalizes the essence of the language mechanism and is shown to enjoy consistency and canonicity. Future work can explore ways to further improve the practicality and expressivity of the language mechanism.

REFERENCES

- Martín Abadi, Luca Cardelli, Pierre-Louis Curien, and Jean-Jacques Lévy. 1989. Explicit substitutions. In *ACM SIGPLAN Symp. on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/96709.96712>
- Harold Abelson and Gerald Jay Sussman. 1996. *Structure and interpretation of computer programs*. The MIT Press.
- Thorsten Altenkirch and Ambrus Kaposi. 2016. Type theory in type theory using quotient inductive types. In *ACM SIGPLAN Symp. on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/2837614.2837638>
- David Aspinall. 1995. Subtyping with singleton types. In *Int'l Workshop on Computer Science Logic*. <https://doi.org/10.1007/BFb0022243>
- Don Batory, Peter Höfner, and Jongwook Kim. 2011. Feature interactions, products, and composition. In *ACM Int'l Conf. on Generative Programming and Component Engineering (GPCE)*. <https://doi.org/10.1145/2047862.2047867>
- Matthias Blume, Umut A. Acar, and Wonseok Chae. 2006. Extensible programming with first-class cases. In *ACM SIGPLAN Conf. on Functional Programming (ICFP)*. <https://doi.org/10.1145/1159803.1159836>
- Olivier Boite. 2004. Proof reuse with extended inductive types. In *Int'l Conf. on Theorem Proving in Higher Order Logics*.
- Gilad Bracha and William Cook. 1990. Mixin-based inheritance. In *ACM SIGPLAN Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*. <https://doi.org/10.1145/97945.97982>
- John Cartmell. 1986. Generalised algebraic theories and contextual categories. *Annals of Pure and Applied Logic* 32 (1986).
- William R. Cook, Walter L. Hill, and Peter S. Canning. 1990. Inheritance is not subtyping. In *ACM SIGPLAN Symp. on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/96709.96721>
- Coq [n.d.]. The Coq proof assistant. <https://coq.inria.fr>.
- Thierry Coquand. 2013. Presheaf model of type theory. (2013). <https://www.cse.chalmers.se/~coquand/presheaf.pdf>
- Thierry Coquand. 2019. Canonicity and normalization for dependent type theory. *Theoretical Computer Science* 777, C (2019). <https://doi.org/10.1016/j.tcs.2019.01.015> arXiv:1810.09367
- Patrick Cousot. 1997. Types as abstract interpretations. In *ACM SIGPLAN Symp. on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/263699.263744>
- Nicolaas Govert de Bruijn. 1972. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *Indagationes Mathematicae* 75, 5 (1972). [https://doi.org/10.1016/1385-7258\(72\)90034-0](https://doi.org/10.1016/1385-7258(72)90034-0)
- Benjamin Delaware, William Cook, and Don Batory. 2011. Product lines of theorems. *ACM SIGPLAN Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*. <https://doi.org/10.1145/2076021.2048113>
- Benjamin Delaware, Bruno C. d. S. Oliveira, and Tom Schrijvers. 2013a. Meta-theory à la carte. In *ACM SIGPLAN Symp. on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/2429069.2429094>
- Benjamin Delaware, Steven Keuchel, Tom Schrijvers, and Bruno C.d.S. Oliveira. 2013b. Modular monadic meta-theory. In *ACM SIGPLAN Conf. on Functional Programming (ICFP)*. <https://doi.org/10.1145/2500365.2500587>
- Peter Dybjer. 1995. Internal type theory. In *Int'l Workshop on Types for Proofs and Programs*.
- Erik Ernst. 2001. Family polymorphism. In *European Conf. on Object-Oriented Programming (ECOOP)*. https://doi.org/10.1007/3-540-45337-7_17
- Erik Ernst, Klaus Ostermann, and William R. Cook. 2006. A virtual class calculus. In *ACM SIGPLAN Symp. on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/1111037.1111062>
- Matthias Felleisen and Daniel P. Friedman. 1986. *Control Operators, the SECD-machine, and the λ -calculus*. Technical Report. Computer Science Department, Indiana University.
- Matthew Flatt, Shriram Krishnamurthi, and Matthias Felleisen. 1998. Classes and mixins. In *ACM SIGPLAN Symp. on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/268946.268961>
- Yannick Forster and Kathrin Stark. 2020. Coq à la carte: A practical approach to modular syntax with binders. In *ACM SIGPLAN Conf. on Certified Programs and Proofs (CPP)*. <https://doi.org/10.1145/3372885.3373817>
- Daniel Gratzer, Jonathan Sterling, and Lars Birkedal. 2019. Implementing a modal dependent type theory. *Proc. of the ACM on Programming Languages (PACMPL)* 3, ICFP (2019). <https://doi.org/10.1145/3341711>
- Jasper Hugunin. 2017. A Coq development of the theory of indexed W types with function extensionality. <https://github.com/jashug/IWTypes>.
- Jasper Hugunin. 2020. Why not W?. In *Int'l Conf. on Types for Proofs and Programs (TYPES)*. <https://doi.org/10.4230/LIPICs.TYPES.2020.8>

- Antonius J. C. Hurkens. 1995. A simplification of Girard’s paradox. In *Int’l Conf. on Typed Lambda Calculi and Applications*.
- Einar Broch Johnsen and Christoph Lüth. 2004. Theorem reuse by proof term transformation. In *Int’l Conf. on Theorem Proving in Higher Order Logics*. https://doi.org/10.1007/978-3-540-30142-4_12
- Ambrus Kaposi. 2017. *Type theory in a type theory with quotient inductive types*. Ph.D. Dissertation. University of Nottingham.
- Ambrus Kaposi, Simon Huber, and Christian Sattler. 2019. Gluing for type theory. In *Int’l Conf. on Formal Structures for Computation and Deduction (FSCD)*. <https://doi.org/10.4230/LIPIcs.FSCD.2019.25>
- Steven Keuchel and Tom Schrijvers. 2013. Generic datatypes à la carte. In *9th ACM SIGPLAN Workshop on Generic Programming*.
- Yao Li and Stephanie Weirich. 2022. Program adverbs and Tlön embeddings. *Proc. of the ACM on Programming Languages (PACML)* 6, ICFP (2022). <https://doi.org/10.1145/3547632> arXiv:2207.05227
- Zhaohui Luo. 2012. Notes on universes in type theory. <http://www.cs.rhul.ac.uk/home/zhaohui/universes.pdf> Notes for a talk at the Institute for Advanced Study in Princeton.
- Ole L. Madsen and Birger Møller-Pedersen. 1989. Virtual classes: A powerful mechanism in object-oriented programming. In *ACM SIGPLAN Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*. <https://doi.org/10.1145/74877.74919>
- Per Martin-Löf. 1982. Constructive mathematics and computer programming. In *Logic, Methodology and Philosophy of Science VI. Studies in Logic and the Foundations of Mathematics*, Vol. 104. [https://doi.org/10.1016/S0049-237X\(09\)70189-2](https://doi.org/10.1016/S0049-237X(09)70189-2)
- Per Martin-Löf. 1984. *Intuitionistic Type Theory: Notes by Giovanni Sambin of a Series of Lectures Given in Padua, June 1980 (Studies in Proof Theory)*.
- Per Martin-Löf. 1992. Substitution calculus. Notes from a lecture given in Göteborg.
- Peter Morris and Thorsten Altenkirch. 2009. Indexed containers. In *Symp. on Logic In Computer Science (LICS)*, Vol. 10.
- Anne Mulhern. 2006. Proof weaving. In *1st Informal ACM SIGPLAN Workshop on Mechanizing Metatheory*.
- Nathaniel Nystrom, Stephen Chong, and Andrew C. Myers. 2004. Scalable extensibility via nested inheritance. In *ACM SIGPLAN Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*. <https://doi.org/10.1145/1028976.1028986>
- Martin Odersky and Matthias Zenger. 2005. Scalable component abstractions. In *ACM SIGPLAN Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*. <https://doi.org/10.1145/1094811.1094815>
- Frank Pfenning. 2009. Lecture notes on harmony.
- Benjamin C. Pierce, Arthur Azevedo de Amorim, Chris Casinghino, Marco Gaboardi, Michael Greenberg, Cătălin Hrițcu, Vilhelm Sjöberg, and Brent Yorgey. 2022. Software foundations: Volume 2 (programming language foundations). (2022). Version 6.2.
- John C. Reynolds. 1975. User-defined types and procedural data structures as complementary approaches to data abstraction. In *New Directions in Algorithmic Languages*. https://doi.org/10.1007/978-1-4612-6315-9_22
- Talia Ringer, RanDair Porter, Nathaniel Yazdani, John Leo, and Dan Grossman. 2021. Proof repair across type equivalences. In *ACM SIGPLAN Conf. on Programming Language Design and Implementation (PLDI)*. <https://doi.org/10.1145/3453483.3454033>
- Andreas Rossberg and Derek Dreyer. 2013. Mixin’ up the ML module system. *ACM Tran. on Programming Languages and Systems (TOPLAS)* 35, 1 (April 2013). <https://doi.org/10.1145/2450136.2450137>
- Christopher Schwaab and Jeremy G. Siek. 2013. Modular type-safety proofs in Agda. In *7th Workshop on Programming Languages Meets Program Verification*. <https://doi.org/10.1145/2428116.2428120>
- Jonathan Sterling. 2019. Algebraic type theory and universe hierarchies. (2019). arXiv:1902.08848
- Christopher A. Stone. 2000. *Singleton kinds and singleton types*. Ph.D. Dissertation. Carnegie Mellon University.
- Christopher A. Stone and Robert Harper. 2000. Deciding type equivalence in a language with singleton kinds. In *ACM SIGPLAN Symp. on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/325694.325724>
- Wouter Swierstra. 2008. Data types à la carte. *Journal of Functional Programming (JFP)* 18, 4 (2008). <https://doi.org/10.1017/S0956796808006758>
- Kresten Krab Thorup. 1997. Genericity in Java with virtual types. In *European Conf. on Object-Oriented Programming (ECOOP)*. <https://doi.org/10.1007/BFb0053390>
- The Univalent Foundations Program. 2013. *Homotopy Type Theory: Univalent Foundations of Mathematics*. <https://homotopytypetheory.org/book>. arXiv:1308.0729

- Cas van der Rest, Casper Bach Poulsen, Arjen Rouvoet, Eelco Visser, and Peter Mosses. 2022. Intrinsically-typed definitional interpreters à la carte. *Proc. of the ACM on Programming Languages (PACMPL)* 6, OOPSLA2 (2022). <https://doi.org/10.1145/3563355>
- Philip Wadler et al. 1998. The expression problem. <http://homepages.inf.ed.ac.uk/wadler/papers/expression/expression.txt>
Discussion on the Java-genericity mailing list.
- Yizhou Zhang and Andrew C. Myers. 2017. Familia: Unifying interfaces, type classes, and family polymorphism. *Proc. of the ACM on Programming Languages (PACMPL)* 1, OOPSLA (Oct. 2017). <https://doi.org/10.1145/3133894>