YIYUN LIU, University of Pennsylvania, USA

STEPHANIE WEIRICH, University of Pennsylvania, USA

In a dependent type theory with β -equivalence as its equational theory, the confluence of untyped reduction and normalization immediately give us a proof of the decidability of type conversion, where the decision procedure for convertibility simply checks the equality of the β -normal forms of its inputs. This technique is not available in the presence of surjective pairing (i.e. the η -law for pairs) because $\beta\eta$ -reduction is not confluent. In this work, we show that by adopting established syntactic techniques, we can resolve the issue with confluence caused by surjective pairing, and recover a confluence-based proof of decidability of type conversion. Compared to existing proof developments, which rely on semantic tools such as Kripke-style logical relations, our proof modularly composes a minimal semantic proof of untyped normalization and a syntactic proof of decidability. This modularity enables us to explore algorithmic conversion through syntactic methods without modifying the minimal semantic proof. We have fully mechanized our results using the Rocq theorem prover.

1 Introduction

A prerequisite to implementing a type checker for a dependent type theory is a decision procedure for type conversion. However, finding such an algorithm is nontrivial; the type theories' specification of definitional equality often does not directly induce an algorithm, due to the rules such as symmetry and transitivity. Furthermore, not only do we want an algorithm, but we also want to know that it is correct, *i.e.* we want to prove that it is sound and complete with respect to definitional equality. By mechanizing such a correctness proof in a proof assistant, we can extract a certified conversion routine and have the highest level of assurance in our implementation.

Designers and implementors of existing type theories take various approaches to this problem. 27 When definitional equality is defined as *untyped* β -equivalence, the confluence of untyped β -28 reduction gives us a reliable formula for both finding algorithms and reasoning about them. To 29 test the β -equivalence of two terms, confluence says it suffices to find a common term that both 30 inputs normalize to through any reduction strategy. As a result, we obtain a class of algorithms 31 that *reduce* their inputs to β -normal forms and then *compare* for α -equivalence. The completeness 32 and soundness of these algorithms follows directly from confluence. To prove that these algorithms 33 terminate, and thus that conversion is decidable, we can use a logical predicate [Girard et al. 1989] 34 to show that all well-typed terms are normalizing under any reduction strategy. 35

We can also use the confluence-based method for type theories that use *typed* β -equivalence for conversion. Adams [2006] and Siles and Herbelin [2012] show the equivalence between pure type systems [Barendregt 1991] that use untyped and typed β -equality. Thus, we can transport decidability results for a system with untyped definitional equality to its equivalent system with typed definitional equality.

However, challenges arise when definitional equality includes of η -laws, especially the pair η -law, also known as *surjective pairing*. While confluence holds for Curry-style systems that include the function η -law, surjective pairing breaks confluence for both Curry and Church-style systems [Barendregt 1993; Geuvers 1993; Klop and de Vrijer 1989; Lennon-Bertrand 2022]. Without confluence, we are unable to justify the correctness of the reduce-and-compare algorithm.

Authors' Contact Information: Yiyun Liu, University of Pennsylvania, Philadelphia, USA, liuyiyun@seas.upenn.edu;
 Stephanie Weirich, University of Pennsylvania, Philadelphia, USA, sweirich@seas.upenn.edu.

1 2

3 4

5

6

7

8

9

10

11

12

13

49

39

40

41

42

43

44

58

59

60

76

Furthermore, η -reduction for pairs is not type preserving [Abel and Coquand 2007], making it 50 difficult to show the equivalence between systems with typed $\beta\eta$ -equality and untyped $\beta\eta$ -equality. 51 Because typed definitional equality only relates well-typed terms, it is more restrictive than untyped 52 equality. In particular, in the transitivity rule, to derive A = C from A = B and B = C, typed equality 53 also requires that the intermediate type B be well-typed. However, without type preservation, it is 54 not obvious how the untyped algorithm that compares $\beta\eta$ -normal forms is sound with respect to a 55 typed definitional equality, as the $\beta\eta$ -reduction sequence to the normal form may contain ill-typed 56 57 terms.

These technical issues are unfortunate because the pair η -law is important for the expressiveness of the type system! In theorem provers, the η -law for pairs generalizes to the η -laws for non-empty dependent records, allowing us to type-check more terms.

In this work, we resolve these issues and show that reduce-and-compare algorithms with η laws for functions and pairs correctly implement type conversion. Our key observation is that $\beta\eta$ -reduction with the function η -law and surjective pairing is confluent on an inductively defined set of strongly normalizing terms (originally defined by Van Raamsdonk and Severi [1995] and denoted as SN). By applying the *SN*-method of Joachimski and Matthes [2003] to formulate our logical predicate, we can prove that all well-typed terms are in the SN set, allowing us to directly show the completeness of reduce-and-compare algorithms.

To show the soundness of any reduce-and-compare algorithm, we adopt the syntactic technique 68 of Goguen [2005] and introduce Coquand's untyped conversion algorithm [Coquand 1991] as an 69 intermediate step. Coquand's algorithm performs η -expansion based on the shape of the terms 70 being compared. The shape-based η -expansion is type-preserving and therefore enables us to give a 71 direct proof of its soundness with respect to the typed convertibility. By proving the completeness 72 of Coquand's algorithm with respect to the untyped reduce-and-compare algorithm, we fully 73 establish the equivalence between the typed convertibility, Coquand's algorithm, and the untyped 74 reduce-and-compare algorithm. 75

Why yet another decidability proof of type conversion? To motivate our confluence-based algo-77 rithms, we start by reviewing the state-of-the-art proof techniques for the decidability of conversion. 78 Abel and Coquand [2007] and later efforts by Abel et al. [2017]; Abel and Scherer [2012] and Adjedj 79 et al. [2024], prove the decidability of type conversion by interpreting types as partial equivalence 80 relations (PER). Instead of modeling the definitional equality with untyped $\beta\eta$ -reduction, this 81 approach models definitional equality using PERs. The extensionality properties of the PER model, 82 justifies the η -laws for functions and products. This approach sidesteps the need for the confluence, 83 so works in situations where untyped $\beta\eta$ -reduction cannot model definitional equality (e.g. the 84 inclusion of η -laws for singleton types). However, the power of the PER model comes at a cost. 85 Regardless of whether the algorithm being justified is typed [Abel and Scherer 2012; Harper and 86 Pfenning 2005] or untyped [Abel and Coquand 2007], the PER method can only be carried out 87 on type systems with a *typed* definitional equality. For systems that do not have an equivalent 88 form with a typed definitional equality, such as ICC [Barras and Bernardo 2008; Miquel 2001], 89 DDC [Choudhury et al. 2022], and DCOI [Liu et al. 2024b], this approach is not applicable. 90

Coquand [2019], Altenkirch and Kaposi [2016], and the more recent synthetic method by Sterling and Angiuli [2021] prove the decidability of type conversion by working with an algebraic structure where syntactic terms are quotiented by definitional equality. The algebraic approach completely eschews the notion of reduction and proves the correctness of the normalization by evaluation (NbE) algorithm [Berger et al. 1998] without reasoning about reduction at all. Similar to the PER approach, the algebraic reduction-free approach does not rely on the confluence of $\beta\eta$ -reduction, but shares the limitation that the object language must have a typed definitional equality. While

99 the reduction-free approach leads to a simple proof on paper, it remains difficult to carry it out in 100 existing intensional theorem provers [Adjedj et al. 2024].

Despite the many advantages of the PER method and the algebraic reduction-free method, we believe the confluence-based proof is appealing for the following reasons. First, we can directly apply the confluence-based method to a system with an untyped definitional equality without taking a detour to typed definitional equality. Avoiding such detours not only leads to simpler proofs, but is also necessary for systems that are not known to be compatible with typed equality.

Second, because the confluence result is untyped, we can reuse it across multiple different type
 systems. The presence of expressive type system features such as large eliminations, subtyping, or
 inductive datatypes is independent of our methodology.

Finally, the confluence-based method leads to a modular proof consisting of a minimal, localized semantic proof of normalization, and a purely syntactic proof of decidability parameterized by confluence and normalization. The clean decoupling between the semantic and syntactic arguments allows us to mechanize more of the decidability proof in a weak metatheory by assuming only normalization [Sozeau et al. 2019].

We have mechanized all of the results of this paper using the Rocq proof assistant [Team 2024] and the proof development is available in the supplementary material under the source subdirectory. By leveraging the method of encoding general recursion by Bove and Capretta [2005], we can extract the verified conversion checker as an executable OCaml program.

The structure of our proof scripts supports our claim about its modularity and reusability. The 118 entire development has approximately 10,000 LoC. The logical predicate, which is used to prove 119 normalization, consists of 1,500 LoC (15% of the total). In contrast, the logical relations found 120 in both Abel et al. [2017] and Wieczorek and Biernacki [2018] take up more than 50% of their 121 respective code bases. Our semantic proofs are therefore significantly more lightweight. Our proof 122 for $\beta\eta$ -confluence and various other untyped injectivity properties, all of which are agnostic to 123 the type system, consists of 5,000 LoC (50%). The remaining proofs are syntactic metatheoretic 124 results of the type theory and syntactic proofs that relate Coquand's algorithm to $\beta\eta$ -reduction. 125 Experimenting with new variations of the conversion algorithm would thus only require changes 126 to the remaining 3,500 lines of syntactic proof that are specific to the algorithm. 127

128 129

130

131

132

133

134

135

136

137

138

139

140

141

142

143

144

145

146 147 *Our contributions.* In this work, we analyze an expressive dependent type theory, called $\lambda^{\Pi,\Sigma,U_i,\mathbb{N}}$, that features large eliminations, a cumulative universe hierarchy with subtyping, and a typed convertibility relation with η -laws for functions and dependent pairs. Section 2 presents the syntax and typing specification of our object language and its basic syntactic results.

In Section 3, we prove the confluence of $\beta\eta$ -reduction for SN and apply the *SN*-method in a dependently typed setting to prove that every well-typed term is in SN using an untyped logical predicate. The confluence result over SN allows us to model convertibility with untyped $\beta\eta$ -reduction and conclude the completeness of the reduce-and-compare algorithm with respect to the definitional subtyping relation.

In Section 4, we introduce Coquand's algorithm for type conversion extended with subtyping and surjective pairing as bridge between the untyped reduce-and-compare algorithm and typed definitional subtyping. Leveraging the fact that the η -expansion performed in Coquand's algorithm is type-preserving, we prove the soundness of Coquand's algorithm with respect to the typed definitional subtyping. We prove the completeness of Coquand's algorithm with respect to untyped $\beta\eta$ -equivalence using a termination metric adapted from Goguen [2005]. We are thus able to conclude our decidability proof by showing that Coquand's algorithm and the reduce-and-compare algorithm (with well-typed inputs) are both equivalent to the typed convertibility relation.

(Context well-formedness)

¹⁴⁸ While our object language, $\lambda^{\Pi, \Sigma, U_i, \mathbb{N}}$, mostly resembles the core languages of intensional type ¹⁴⁹ theories, there are a few subtleties in its design, which we discuss in Section 5. In particular, we ¹⁵⁰ include type constructor injectivity rules as part of the subtyping relation (Section 5.1), and use ¹⁵¹ of Curry-style lambda terms (Section 5.2). We also deepen our comparison to related systems in ¹⁵² Section 6 and conclude in Section 7.

Typing contexts

variables, type universes

function types, abstractions, applications

Terms

WF-CONS

 $\frac{\vdash \Gamma \qquad \Gamma \vdash A : \mathcal{U}_i}{\vdash \Gamma, x : A}$

 $|\Sigma x: A. B| (a, b) | \pi_1 x | \pi_2 x$ pair types, pairs, projections

Fig. 1. Grammar, excluding natural numbers

2 Specification of $\lambda^{\Pi,\Sigma,U_i,\mathbb{N}}$

A, B, C, a, b, c, P ::=

 $\Gamma ::= \cdot | \Gamma, x : A$

 $|\Pi x: A. B | \lambda x. b | b a$

WF-NIL

⊢ .

 $|x|\mathcal{U}_i$

1	5	6
1	5	7
1	5	8
1	5	9

153

154 155 4

159 160

161

162

163

164

165 166

167

169

170

171

172

173

1

1 1 1

168

⊢ Γ

 $\Gamma \vdash a : A$ (Typing) WT-PI WT-APP $\Gamma \vdash A : \mathcal{U}_i$ $\Gamma \vdash b : \Pi x : A. B$ WT-VAR Wт-Abs $\Gamma, x : A \vdash B : \mathcal{U}_j$ $\Gamma, x : A \vdash b : B$ **⊢**Γ $x : A \in \Gamma$ $\Gamma \vdash a : A$ $\frac{1}{\Gamma \vdash \Pi \mathbf{r} \cdot A \cdot B \cdot \mathcal{I}_{I_{\text{EVE}}}}$ $\frac{1}{\Gamma \vdash \lambda r \ h \cdot \Pi r \cdot A \ B}$ $\Gamma \vdash ha \cdot B[a/r]$ $\Gamma \vdash r \cdot A$

74	1 7 7 7 11		<i>ivj</i> i <i>i i i i i i i i i i</i>	$\mathbf{I} + \mathbf{v} \mathbf{u} + \mathbf{D} [\mathbf{u} / \mathbf{x}]$
75		Wt-Sig	Wt-Pair	
76	WT-UNIV	$\Gamma \vdash A : \mathcal{U}_i$	$\Gamma \vdash a : A \qquad \Gamma \vdash b : B[a/x]$	WT-Proj1
77	$\vdash \Gamma$	$\Gamma, x : A \vdash B : \mathcal{U}_j$	$\Gamma \vdash \Sigma x : A. B : \mathcal{U}_i$	$\Gamma \vdash a : \Sigma x : A. B$
78 79	$\overline{\Gamma \vdash \mathcal{U}_i : \mathcal{U}_{i+1}}$	$\Gamma \vdash \Sigma x : A. B : \mathcal{U}_{i \lor j}$	$\Gamma \vdash (a, b) : \Sigma x : A. B$	$\Gamma \vdash \pi_1 a : A$
80		Wt-Proj2	Wt-Conv	
81		$\Gamma \vdash a : \Sigma x : A. B$	$\Gamma \vdash a : A \qquad \Gamma \vdash A \le B$	}
82				-
83		$\Gamma \vdash \pi_2 a : B[\pi_1 a/x]$	$\Gamma \vdash a : B$	

182 183 184

185

186

187

188

Fig. 2. Typing rules

We start by presenting $\lambda^{\Pi,\Sigma,U_i,\mathbb{N}}$, our dependently typed object language, including its grammar and typing rules (Figures 1 to 3). These rules specify *what* programs type check, but they do not directly describe a type checking algorithm.

The language includes functions ($\lambda x. a$), dependent pairs (a, b) with projection operators ($\pi_1 a$ and $\pi_2 a$), and a Russell-style infinite universe hierarchy (\mathcal{U}_i). Function abstractions do not include type annotations; we discuss in Section 5.2 our proofs can be easily adapted to a Church-style system.

The language in our mechanized proofs also includes natural numbers with an induction principle, but we omit these forms from the text. We include natural numbers in our development to provide an observable infinite data type so our language supports large eliminations. Furthermore, natural

$F \vdash a = b : A$			(Equality	
		E-Trans	E-AbsExt	
E-Refl	E-Sym	$\Gamma \vdash a = b : A$	$x \notin \Gamma$	
$\Gamma \vdash a : A$	$\Gamma \vdash b = a : A$	$\Gamma \vdash b = c : A$	$\Gamma, x : A \vdash a x = b x : B$	
$\Gamma \vdash a = a : A$	$\Gamma \vdash a = b : A$	$\Gamma \vdash a = c : A$	$\Gamma \vdash a = b : \Pi x : A. B$	
E-PairExt				
$\Gamma \vdash \pi$	$a = \pi_1 b : A$	E-AppAbs		
$\Gamma \vdash \pi_2 a =$	$= \pi_2 b : B[\pi_1 a/x]$	$\Gamma \vdash b : A$	$\Gamma, x : A \vdash a : B$	
$\Gamma \vdash a$	$= b : \Sigma x : A. B$	$\overline{\Gamma \vdash (\lambda x. a) \ b}$	= a[b/x] : B[b/x]	
E-ProjPai	r1	E-ProjPair:	2	
$\Gamma \vdash a : A$	$\Gamma \vdash b : B[a/x]$	$\Gamma \vdash a : A$	$\Gamma \vdash b : B[a/x]$	
Γ + <i>ι</i>	$\tau_1(a,b)=a:A$	$\Gamma \vdash \pi_2(a, b) = b : B[a/x]$		
$\Gamma \vdash A \leq B$			(Subtyping	
			L-PI	
-Trans	L-Eq	L-Univ	$\Gamma \vdash A_1 \le A_0$	
$\vdash A \leq B \qquad \Gamma \vdash B \leq$	$C \qquad \Gamma \vdash A = B : \mathcal{U}_i$	$\vdash \Gamma \qquad i \leq j$	$\Gamma, x : A_0 \vdash B_0 \le B_1$	
$\Gamma \vdash A \leq C$	$\Gamma \vdash A \leq B$	$\overline{\Gamma \vdash \mathcal{U}_i \leq \mathcal{U}_j}$	$\overline{\Gamma \vdash \Pi x : A_0. B_0} \le \Pi x : A_1. B_1$	
L-Sig			L-PiProj2	
$\Gamma \vdash A_0 \leq A_1$	L-PiProj1		$\Gamma \vdash \Pi x : A_0. B_0 \leq \Pi x : A_1. B_1$	
$\Gamma, x : A_0 \vdash B_0 \leq B$	$\Gamma \vdash \Pi x : A_0.$	$B_0 \leq \Pi x : A_1 . B_1$	$\Gamma \vdash a_0 = a_1 : A_1$	
$\Gamma \vdash \Sigma x : A_0. B_0 \le \Sigma x : A_0$	$A_1. B_1$ $\Gamma \vdash $	$A_1 \le A_0$	$\Gamma \vdash B_0[a_0/x] \le B_1[a_1/x]$	

Fig. 3. Equality and subtyping rules (selected)

numbers help us ensure that our proof techniques are robust under the addition of positive types (i.e. types with pattern-matching as their elimination forms). This addition leads to no surprises and requires no special treatment during our development. The full set of rules that includes natural numbers can be found in the supplementary material.

The $\lambda^{\Pi,\Sigma,U_i,\mathbb{N}}$ type system includes subtyping. Rule WT-CONV refers to the definitional subtyping relation $\Gamma \vdash A \leq B$, which is defined in Figure 3 along with the $\beta\eta$ -rules for the definitional equality $\Gamma \vdash a = b : A$. The rules for equality are standard, except for the addition of the injectivity rules L-PIPROJONE and L-PIPROJTWO and the omitted counterpart for Σ types. These rules simplify our proofs and do not affect decidability; we discuss them further in Section 5.1.

This type system satisfies standard syntactic properties, such as weakening and substitution (not shown). In addition, our mechanization includes proofs of the following. We can recover the well-formedness of context from the mutually defined typing and conversion judgments.

LEMMA 2.1 (CONTEXT REGULARITY). If
$$\Gamma \vdash a : A, \Gamma \vdash a = b : A, or \Gamma \vdash A \leq B, then \vdash \Gamma$$
.

In the presence of rule WT-CONV, the generation lemma is useful for recovering information about a typing derivation.

LEMMA 2.2 (GENERATION (SELECTED)).

- If $\Gamma \vdash \lambda x$. a : C, then there exists some A and B such that $\Gamma \vdash \Pi x : A$. $B \leq C$ and $\Gamma, x : A \vdash a : B$.

- If $\Gamma \vdash ba : C$, then there exists some A and B such that $\Gamma \vdash b : \Pi x : A. B, \Gamma \vdash a : A$ and 246 $\Gamma \vdash B[a/x] \leq C.$
 - If $\Gamma \vdash \Pi x : A, B : C$, then there exists some i such that $\Gamma \vdash A : \mathcal{U}_i, \Gamma, x : A \vdash B : \mathcal{U}_i$, and $\Gamma \vdash \Pi x : A. B : \mathcal{U}_i \text{ and } \Gamma \vdash \mathcal{U}_i \leq C.$
 - If $\Gamma \vdash \mathcal{U}_i : C$, then $\Gamma \vdash \mathcal{U}_i : \mathcal{U}_{i+1}$ and $\Gamma \vdash \mathcal{U}_{i+1} \leq C$.

Using the substitution lemma and rules L-PIPROJONE and L-PIPROJTWO, we can prove that β -reduction preserves typing. We defer the precise definition of β -reduction to the next section.

LEMMA 2.3 (SUBJECT REDUCTION). If $\Gamma \vdash a : A$ and $a \rightsquigarrow_{\beta} b$, then $\Gamma \vdash b : B$.

Finally, regularity states that if a term is well-typed, then its type must be well-formed.

LEMMA 2.4 (REGULARITY).

- If $\Gamma \vdash a : A$, then there exists i such that $\Gamma \vdash A : \mathcal{U}_i$.
- If $\Gamma \vdash a = b : A$, then $\Gamma \vdash a : A$ and $\Gamma \vdash b : A$.
- If $\Gamma \vdash A \leq B$, then there exists i such that $\Gamma \vdash A : \mathcal{U}_i$ and $\Gamma \vdash B : \mathcal{U}_i$.

Confluence of $\beta\eta$ **-Reduction for Strong Normalizing Terms** 3

Our goal in this section is to show that we can implement the subtyping relation $\Gamma \vdash A \leq B$. A key step in this process is to model the typed convertibility relation $\Gamma \vdash a = b : A$ with an untyped reduce-and-compare algorithm defined in terms of $\beta\eta$ -reduction. This algorithm, which we introduce in Section 3.6, is not transitive by definition, a necessary property if we are to use it to model the transitivity rule of typed convertibility. Instead, we use the confluence of $\beta\eta$ -reduction to show that our algorithmic relation is transitive.

However, the usual $\beta\eta$ -reduction relation is not confluent, due to surjective pairing. (See Klop [1980] for a counterexample). We solve this issue by identifying a set of terms where confluence holds-an inductively characterized set, called SN, that contains only strongly normalizing terms [Van Raamsdonk and Severi 1995]-and then arguing that all terms that we care about are in this set. Following Joachimski and Matthes [2003], we use an untyped logical predicate to show that all well-typed terms are in SN. Because we only invoke the conversion algorithm on well-typed terms, the confluence result for SN is sufficient to justify the transitivity of the reduce-and-compare algorithm.

3.1 **Reduction Relations and Normal Forms**



Before we introduce SN, we start by reviewing standard definitions that help us both motivate and state properties about this set.

292

293 294

262

263

264

265

266

267

268

269

270

271

272

273

274

275

276

277 278 6

247

248

249

250 251

252

253

254 255

First, in Figure 4, we define primitive β and η -reductions. Rule PE-PAIRETA is also known as surjective pairing. We refer to a lambda term that can be reduced by one of these primitive reduction rules as a β or η -redex respectively.

We use \rightarrow to denote the strong/full reduction relation, which can reduce anywhere in the term, including under constructor forms, and \Rightarrow to denote its parallel variant [Barendregt 1993; Takahashi 1995]. To be specific about which primitive rules may participate in full reduction, we use subscripts. Thus, the relations \rightarrow_{β} , \rightarrow_{η} , and $\rightarrow_{\beta\eta}$ denote the full β , η , and $\beta\eta$ -reduction relations. We also use \sim_{lo} to denote the leftmost-outermost β -reduction strategy and \sim_{h} as its corresponding weak reduction variant, which does not reduce under constructors. Note that our proofs do not require versions of these relations that include η -reduction, so we omit the β -subscript.

306	Neutral forms (ne)	ν	::=	$x \mid v u \mid \pi_1 v \mid \pi_2 v$
307	Normal forms (nf)	и	::=	$\Pi x : u. u \mid \Sigma x : u. u \mid \lambda x. u \mid (u, u) \mid \mathcal{U}_i$
308				
309	Weak-head neutral forms (whne)	е	::=	$x \mid e \mid \pi_1 e \mid \pi_2 e$
310	Canonical forms (canf)	h	::=	$\Pi x: A. B \mid \Sigma x: A. B \mid \lambda x. a \mid (a, b) \mid \mathcal{U}_i$
311	Weak-head normal forms (whnf)	f	::=	$e \mid h$
312	Fig. F. Crowner			and normal formers

Fig. 5. Grammars for neutral and normal forms

Figure 5 defines β -neutral (ν) and β -normal (u) forms and their weak-head counterparts. An expression that is in weak-head normal form but is not neutral is called *canonical*. It is straightforward to verify that neither neutral nor normal forms contain any β -redexes and that the η -reduction relation preserves β -normal forms.

LEMMA 3.1. If $a \in \mathbf{nf}$ or $a \in \mathbf{ne}$, then $a \nleftrightarrow_{\beta} b$.

LEMMA 3.2. If $u \rightsquigarrow_{\eta} a$, then $a \in \mathbf{nf}$. If $v \rightsquigarrow_{\eta} a$, then $a \in \mathbf{ne}$.

Note that the converse of Lemma 3.1 is not true—the inductive definitions of neutral and normal forms rule out stuck terms. For examples, the terms $\pi_1 (\lambda x. x)$ or (x, y) z are unable to take any further β -steps but are *not* in normal form as defined by Figure 5.

3.2 An Inductive Characterization of Strongly Normalizing Terms

Following Van Raamsdonk and Severi [1995], we define the set SN of strongly normalizing terms mutually with the set of strongly neutral terms SNe and the relation $a \sim_{SN} b$, a subrelation of the weak-head reduction relation $a \sim_h b$. Ignoring N-Exp, the sets SN and SNe are exactly the sets of normal and neutral forms from Figure 5. By adding N-Exp, we can expand the sets to include terms that contain β -redexes but nevertheless reduce to a normal or neutral forms.

In rule N-APPABS, the fact that a[b/x] is strongly normalizing alone does not necessarily imply that (λx . *a*) *b* is also strongly normalizing. We can construct a counterexample by picking *b* to be an infinite loop and *a* an expression that does not contain the variable *x*. Thus, the $b \in SN$ constraint in N-APPABS ensures that if the term *b* is strongly normalizing and $a \rightsquigarrow_{SN} b$, then *a* must also be strongly normalizing.

By forgetting the SN and SNe premises in \sim_{SN} , we can recover from each of these judgments a leftmost-outermost reduction (\sim_{lo}) sequence to normal or neutral forms.

339 Lemma 3.3 (SN leftmost-outermost reduction).

- If $a \in SN$, then there exists u such that $a \sim_{lo}^{*} u$.
- If $a \in SNe$, then there exists v such that $a \sim_{lo}^{*} v$.
- If $a \rightsquigarrow_{SN} b$, then $a \rightsquigarrow_h b$.

343

305

313

318

319

320 321

322

323

324 325

326

327

328

329

330

331

332

333

334

335

336

337

338

340

N-PAIR $a \in SN$		N-Pt			
	$b \in SN$	$A \in SN$	$B \in \mathbf{SN}$	$A \in \mathbf{SN}$	$B \in \mathbf{SN}$
(<i>a</i> , <i>b</i>)	∈ SN	$\Pi x: A$	$A. B \in SN$	$\Sigma x: A.$	$B \in \mathbf{SN}$
NIV	N-SN $a \in S$	e SNe	$\substack{\text{N-Exp}\\ a \rightsquigarrow_{\text{SN}} b}$	$b \in SN$	
∈ SN	$a \in \mathcal{C}$	SN	$a \in S$	SN	
				(Strong N	leutral Forms)
N-Ap $a \in S$	SNe $b \in$	SN	N-Proj1 $a \in SNe$	N-Proj $b \in S$	2 5 Ne
	$a b \in SNe$		$\overline{\pi_1 \ a \in \mathbf{SNe}}$	$\overline{\pi_2 \ a \in}$	SNe
			(Stro	ong Weak He	ad Reduction)
N- b	AppCong $\in SN$ a_0	$\sim_{\rm SN} a_1$	N-ProjPair1 $b \in SN$	N-Pi	$a \in SN$
$\left[\frac{1}{2} \right]$	$a_0 b \sim_{SN}$	$a_1 b$	$\overline{\pi_1\left(a,b\right)} \rightsquigarrow_{\mathrm{SN}}$	\overline{a} $\overline{\pi_2}$ ($(a, b) \sim_{SN} b$
N-Proj $a \sim$	Cong1 → _{SN} b		N-ProjCong2 <i>a</i> → _{SN} b		
$\overline{\pi_1 a \sim}$	$\rightarrow_{\rm SN} \pi_1 b$		$\overline{\pi_2 \ a \rightsquigarrow_{\mathrm{SN}} \pi_2 \ b}$		
Г -	(a, b) (a, b) $($	$(a, b) \in SN$ $(a, b) = SN$	$(a, b) \in SN$ $\Pi x: A$ $(a, b) \in SN$ $\Pi x: A$ $(a, b) \in SN$ $(a, b) \in SN$ $(a, b) \in SN$ $(a, b) \in SN$ $(a \in SNe)$	Image: Niveral conditionImage: Niveral conditionImage: Niveral conditionNiveral Niveral conditionNiveral Niveral conditionNIV $a \in SN$ $a \in SNe$ $a \in SNe$ $a \sim SN b$ $a \sim SN b$ $a \in SN$ $a \in SNe$ $b \in SN$ $a \in SNe$ $a \in SNe$ $a \in SNe$ $b \in SNe$ $a \in SNe$ $a \in SNe$ $a \in SNe$ $b \in SN$ $a_0 \sim SN a_1$ $b \in SN$ $a = b \in SN$ $a_0 \to SN a_1 b$ $b \in SN$ $a \sim SN b$ $a \sim SN b$ $a \sim SN b$ $a \to SN b$ $a \to SN \pi_1 b$ $a \to SN \pi_2 b$	$ \begin{array}{c c} \hline (a,b) \in \mathrm{SN} & \hline \Pi x: A. \ B \in \mathrm{SN} & \hline \Sigma x: A. \\ \hline \\ \hline \mathrm{NIV} \\ \hline \in \mathrm{SN} & \begin{array}{c} \frac{a \in \mathrm{SNe}}{a \in \mathrm{SN}} & & \\ \hline \\ \frac{a \in \mathrm{SNe}}{a \in \mathrm{SN}} & & \begin{array}{c} \frac{a \sim \mathrm{SN} \ b}{a \in \mathrm{SN}} & \\ \hline \\ \frac{a \in \mathrm{SNe}}{a \in \mathrm{SNe}} & & \begin{array}{c} \frac{a \sim \mathrm{SN} \ b}{a \in \mathrm{SN}} & \\ \hline \\ \frac{a \in \mathrm{SNe}}{a \in \mathrm{SNe}} & & \begin{array}{c} \frac{a \sim \mathrm{SN} \ b}{a \in \mathrm{SN}} & \\ \hline \\ \frac{a \in \mathrm{SNe}}{a \otimes \mathrm{SNe}} & & \begin{array}{c} \frac{a - \mathrm{SNe} \ b}{a \in \mathrm{SN}} & \\ \hline \\ \frac{a \in \mathrm{SNe}}{a \otimes \mathrm{SNe}} & & \begin{array}{c} \frac{a - \mathrm{SNe} \ b}{a \in \mathrm{SN}} & \\ \hline \\ \frac{a \in \mathrm{SNe}}{a \otimes \mathrm{SNe}} & & \begin{array}{c} \frac{a \in \mathrm{SNe} \ b}{a \in \mathrm{SNe}} & \\ \hline \\ \frac{a \in \mathrm{SNe} \ b}{\pi_1 \ a \in \mathrm{SNe}} & & \begin{array}{c} \frac{b \in \mathrm{SN} \ b}{\pi_2 \ a \in \mathrm{SNe}} & \\ \hline \\ \frac{b \in \mathrm{SN} \ a_0 \ b}{\pi_2 \ a \sim \mathrm{SN} \ a_1 \ b} & \\ \end{array} \end{array} $

In the following, we prove that $\beta\eta$ -reduction is confluent by using η -postponement [Barendregt 1993; Takahashi 1995]. A sequence of $\beta\eta$ -reductions starting from an SN term can be factorized into a sequence of β -reductions followed by η -reductions.

An alternative proof of confluence is to use Newman's lemma with the fact that $\beta\eta$ -reduction is locally confluent. A perhaps surprising result, shown by Joachimski and Matthes [2003], is that despite the absence of η -rules in the definition of \sim_{SN} , the set SN also characterizes the set of strongly $\beta\eta$ -normalizing terms. However, the proof that terms in SN are strongly $\beta\eta$ -normalizing, as required by Newman's lemma, is tedious. It requires nested induction and repeated analysis of possible $\beta\eta$ -redexes to show that the set of strongly $\beta\eta$ -normalizing terms satisfies the same congruence rules as SN and SNe.

Our proof based on η -postponement avoids the headache of connecting SN to strong $\beta\eta$ normalization (defined in terms of the well-foundedness of $\sim_{\beta n}$). Furthermore, this proof is also more flexible as it allows us to generalize the $\beta\eta$ -confluence result to weakly β -normalizing systems as we show in Section 5.3.

3.3 **Properties of SN**

Next, we show structural properties of SN to prepare for our η -postponement and confluence proof. Like the inductively defined normal and neutral forms (Figure 5), the sets SN and SNe do not contain stuck terms.

LEMMA 3.4 (NO STUCK TERMS). There is no terms a, b, and c such that $(a, b) c \in SN$, $\pi_1(\lambda x, a) \in SN$, or π_2 ($\lambda x. a$) \in SN.

PROOF. Immediate by case analysis over the derivation of SN.

The inductive characterization of strong normalization satisfies the structural properties of renaming and inversion (not shown), as well as the antisubstitution property below.

Lemma 3.5 (SN antisubstitution).

397

398

399

400

401

407

408

409

410

411

• If $a[b/x] \in SN$, then $a \in SN$.

• If $a[b/x] \in SNe$, then $a \in SNe$.

• If $a[c/x] \rightarrow_{SN} b$, then either $a \in SNe$ or there exists b_0 such that $a \rightarrow_{SN} b_0$ and $b = b_0[c/x]$.

The structural properties are proven by structural induction over the respective derivation except for the application case of inversion that requires Lemma 3.5.

Strong normalization is preserved by both β -reduction and η -reduction. However, because SN and SNe are mutually defined with the reduction relation $a \rightsquigarrow_{SN} b$, we must simultaneously show that $\beta\eta$ -reduction commutes with $a \rightsquigarrow_{SN} b$. This proof of commutativity requires us to generalize the lemma statement to use parallel $\beta\eta$ -reduction. Furthermore, we use $a \rightsquigarrow_{\overline{SN}} b$ to denote the reflexive closure of $a \rightsquigarrow_{SN} b$.

LEMMA 3.6 (SN PRESERVATION).

- If $a \in SN$ and $a \Rightarrow_{\beta\eta} b$, then $b \in SN$.
- If $a \in SNe$ and $a \Rightarrow_{\beta\eta} b$, then $b \in SNe$.
- If $a \sim_{SN} b_0$ and $a \Rightarrow_{\beta\eta} b_1$, then there exists some c such that $b_0 \Rightarrow_{\beta\eta} c$ and $b_1 \sim_{\overline{SN}} c$.

The proof for commutativity (the third case of Lemma 3.6) only works because $a \sim_{SN} b$ is a subrelation of head reduction, which never reduces under constructors. The same commutativity property does not hold if we replace $a \sim_{SN} b$ with full $\beta\eta$ -reduction.

⁴¹⁵ 3.4 Postponement of η -Reduction for SN

⁴¹⁷ While η -postponement is a classic result for the pure untyped lambda calculus, it does not hold ⁴¹⁸ when we add another constructor (such as pairing). A simple counterexample is the expression ⁴¹⁹ $\pi_1 (\lambda x. (y, y) x)$, which normalizes to y by first η -contracting the lambda term $\lambda x. (y, y) x$ to (y, y)⁴²⁰ then β -reducing $\pi_1 (y, y)$. More generally, the counterexamples involve a stuck term that becomes ⁴²¹ unstuck by one step of η -reduction. However, because SN contains no stuck terms, we can prove ⁴²² η -postponement for terms in SN without running into the problematic cases!

The η -postponement theorem that we prove is more precise than the corresponding theorem found in Barendregt [1993] and Takahashi [1995]. We show that every sequence of $\beta\eta$ -reductions can be converted into a sequence of β -reductions followed by a sequence of η -reductions such that *the* η -*reduction sequence does not reduce any* β -*redexes.* We denote this restrictive version of η -reduction as $a \Rightarrow_{ne} b$, mutually defined with the helper reduction relation $a \Rightarrow_{e} b$ in Figure 6. Both relations are subrelations of parallel η -reduction.

These definitions track whether we are reducing the scrutinee of an elimination form. The relation $a \Rightarrow_{ne} b$ can be used when the term *a* is not used in this way. Thus, we can freely perform η -contractions through APPETA and PAIRETA without erasing any β -redex. The $a \Rightarrow_e b$ relation, on the other hand, is used for subterms that are being eliminated and therefore is not allowed to perform any η -reduction steps at the top-level since doing so would erase potential β -redexes.

The desired property that $a \Rightarrow_{ne} b$ does not reduce a β -redex can be stated as follows.

```
<sup>435</sup> LEMMA 3.7 (RESTRICTIVE-\eta AND NF). If a \Rightarrow_{ne} b or a \Rightarrow_{e} b, then b \in \mathbf{nf} implies a \in \mathbf{nf}.
<sup>436</sup>
```

PROOF. Immediate by mutual induction over the derivations of $a \Rightarrow_{ne} b$ and $a \Rightarrow_{e} b$.

⁴³⁸ The following lemma allows us to decompose one step of parallel η -contraction into a sequence ⁴³⁹ of β -reductions followed by one step of restrictive η -contraction. This property requires the SN ⁴⁴⁰ restriction.

441

434

437

Yiyun Liu and Stephanie Weirich

$a \Rightarrow_{ne} b$	(Restrictive η-r	eduction (not eliminated))
EP-AppEta		
$a_0 \notin \operatorname{cant} a_0 \Rightarrow_e a_1$	EP-PairEta	EP-Embed
$x \notin \mathbf{freevar}(a)$	$a_0 \notin \operatorname{canf} \qquad a_0 \Rightarrow_e a_1$	$a \Rightarrow_e b$
$(\lambda x. a_0 x) \Rightarrow_{ne} a_1$	$(\pi_1 a_0, \pi_2 a_0) \Rightarrow_{ne} a_1$	$a \Rightarrow_{ne} \overline{b}$

(eliminated)

NEP-AbsCong	NEP-AppCong	NEP-PairCong
$a_0 \Rightarrow_{ne} a_1$	$a_0 \Rightarrow_e a_1 \qquad b_0 \Rightarrow_{ne} b_1$	$a_0 \Rightarrow_{ne} a_1 \qquad b_0 \Rightarrow_{ne} b_1$
$\overline{\lambda x. a_0} \Rightarrow_e \lambda x. a_1$	$a_0 \ b_0 \Rightarrow_e a_1 \ b_1$	$(a_0, b_0) \Rightarrow_e (a_1, b_1)$

Fig. 6. Definitions of restrictive η -reductions (selected)

LEMMA 3.8 (η -Decomposition). If $a \Rightarrow_{\eta} b$ and $a \in SN$, then there exists c such that $a \rightsquigarrow_{\beta}^{*} c$ and $c \Rightarrow_{ne} b.$

PROOF. By inducting over the derivation of $a \Rightarrow_{\eta} b$. Assuming $a \in SN$, Lemma 3.6 and inversion ensure that every subterm of a or reduced term also belongs to SN. The only interesting cases are EP-APPETA and EP-PAIRETA, which are similar.

We show the EP-APPETA case to demonstrate how the lemmas defined earlier are used. Suppose $a = \lambda x$. $a_0 x$ where $x \notin$ freevar(a) and $a_0 \Rightarrow_{\eta} b$. By induction, there exists some c_0 such that $a_0 \sim^*_{\beta} c_0$ and $c_0 \Rightarrow_{ne} b$. Our goal is to find some c such that $\lambda x. a_0 x \sim^*_{\beta} c$ and $c \Rightarrow_{ne} b$.

From $a_0 \sim \beta_{\beta}^* c_0$, we deduce that $\lambda x. a_0 x \sim \beta_{\beta}^* \lambda x. c_0 x$. When c_0 is not in canonical form, then $a_0 \notin \text{canf}$ holds, and we are allowed to use EP-APPETA to finish off the proof by picking $c = \lambda x. c_0 x.$ The precondition of EP-APPETA requires $c_0 \Rightarrow_e b$ rather than $c_0 \Rightarrow_{ne} b$, but we know that $a_0 \notin can f$.

When c_0 is a canonical form, we know that it must take the form λx . c_1 for some term c_1 , since otherwise Lemmas 3.4 and 3.6 derive a contradiction. We can then pick $c = \lambda x. c_1$ and construct the reduction sequence $a = \lambda x$. $a_0 x \sim_{\beta}^* \lambda x$. $(\lambda x. c_1) x \sim_{\beta} \lambda x. c_1 = c_0 \Rightarrow_{ne} b$.

The parallel η -rules for functions and pairs are hard to analyze as they may reduce any finite number of outermost η -redexes. Lemma 3.8 addresses this problem by decomposing a single step of parallel η -reduction into a sequence of β -reduction followed by one step of \Rightarrow_{ne} , which can only reduce at most one outermost η -redex. Thus, Lemma 3.8 subsumes the notion of k-fold η -expansion from Takahashi [1995] as a tool to make inverting on the derivation of parallel η -reduction tractable.

LEMMA 3.9 (η -postponement). If $a \in SN$, $a \Rightarrow_{\eta} b$, and $b \rightsquigarrow_{\beta} c$, then there exists some b_0 such that $a \rightsquigarrow^*_\beta b_0$ and $b_0 \Rightarrow_\eta c$.

PROOF. By induction on the derivation of $a \Rightarrow_{\eta} b$. Most cases are trivial except for the congruence rules for elimination forms, where Lemma 3.8 is required. We only consider the case for rule P-AppCong as the other cases are similar. Suppose $a = a_0 a_1$ and $b = b_0 b_1$ such that $a_0 \Rightarrow_n b_0$ and $a_1 \Rightarrow_{\eta} b_1$. We invert on the derivation of $b = b_0 b_1 \sim_{\beta} c$ and consider two possible cases. In the first case, $b_0 b_1$ steps into c through one of the congruence rules for application. The goal is then immediate by induction.

In the second case, we have $b_0 b_1 = (\lambda x, b_2) b_1 \sim_{\beta} b_2 [b_1/x] = c$ for some term b_2 . Applying 486 Lemma 3.8 to the hypothesis that $a_0 \Rightarrow_{\eta} b_0 = \lambda x$. b_2 , we deduce that there exists some term a_2 such 487 that $a_0 \rightarrow_{\beta}^* a_2$ and $a_2 \Rightarrow_{ne} \lambda x. b_2$. The proof can be finished off by inverting over the derivation of 488 $a_2 \Rightarrow_{ne} \lambda x. b_2$. Only EP-AppEta and EP-AbsCong could have been used to derive $a_2 \Rightarrow_{ne} \lambda x. b_2$. 489

490

10

 $a \rightarrow$

 $a \Rightarrow_e b$

h

460

461

462

463 464

465

466

467

468

469

470 471

472

473

474

475

476

477

478 479

480

481

482

483

484

485

442

since η -contracting a pair into a lambda term through rule EP-PAIRETA implies the existence of a stuck term that does not belong to SN (Lemma 3.4).

COROLLARY 3.1 (STRENGTHENED η -POSTPONEMENT). If $a \Rightarrow_{\eta} b$ and $b \rightsquigarrow_{\beta} c$, then there exists some b_0 such that $a \rightsquigarrow_{\beta}^* b_0$ and $b_0 \Rightarrow_{ne} c$.

PROOF. Immediate by composing Lemmas 3.8 and 3.9.

By composing Lemma 3.7 and Corollary 3.1, we obtain the final version of our theorem.

COROLLARY 3.2 (η -postponement for normal forms). If $a \sim^*_{\beta\eta} u$, then there exists some u_0 such that $a \sim^*_{\beta} u_0$ and $u_0 \sim^*_{\eta} u$.

Corollary 3.2 states that if some term $\beta\eta$ -reduces to some β -normal form, then we can postpone it in a way such that the intermediate term u_0 remains in β -normal form. The fact that u_0 is in normal form is crucial for deriving the confluence result.

3.5 Confluence of $\beta\eta$ -Reduction for SN

Now, we put everything together and prove confluence using η -postponement. We start by stating the confluence properties of β and η -reductions on their own.

LEMMA 3.10 (CONFLUENCE FOR β). If $a \sim^*_{\beta} b_0$ and $a \sim^*_{\beta} b_1$, then there exists some c such that $b_0 \sim^*_{\beta} c$ and $b_1 \sim^*_{\beta} c$.

LEMMA 3.11 (CONFLUENCE FOR η). If $a \sim_{\eta}^{*} b_0$ and $a \sim_{\eta}^{*} b_1$, then there exists some c such that $b_0 \sim_{\eta}^{*} c$ and $b_1 \sim_{\eta}^{*} c$.

We omit the details of their proofs as they are standard. For our mechanization, we prove confluence for β -reduction using Takahashi's complete development [Takahashi 1995], and confluence for η -reduction through local confluence and Newman's lemma.

Now we can prove our main result.

THEOREM 3.1 (CONFLUENCE OF $\beta\eta$ -REDUCTION FOR SN TERMS). If $a \in SN$, $a \sim_{\beta\eta}^{*} b_0$, and $a \sim_{\beta\eta}^{*} b_1$, then there exists some c such that $b_0 \sim_{\beta\eta}^{*} c$ and $b_1 \sim_{\beta\eta}^{*} c$.

PROOF. By Lemma 3.6, we know that $b_0 \in SN$ and $b_1 \in SN$. Therefore, there must exist some u_0 and u_1 such that $b_0 \sim_{\beta}^* u_0$ and $b_1 \sim_{\beta}^* u_1$. By Lemma 3.11, it suffices to show that u_0 and u_1 are η equivalent. By Corollary 3.2, we can factorize the reduction sequence $a \sim_{\beta\eta}^* u_0$ into $a \sim_{\beta}^* u_2 \sim_{\eta}^* u_0$ for some u_2 . Likewise, we can factorize the sequence $a \sim_{\beta\eta}^* u_1$ into $a \sim_{\beta}^* u_3 \sim_{\eta}^* u_1$ for some u_3 . However, since both u_2 and u_3 are β -normal forms of a, they must be equal according to Lemma 3.10, the confluence of β -reduction. From $u_2 = u_3$, We can immediately conclude that u_0 and u_1 are η -equivalent. The result then follows from the confluence of η -reduction (Lemma 3.11).

⁵³⁰ 3.6 Reduce-and-Compare Algorithms for Equality and Subtyping

Next, we define our reduce-and-compare algorithms. Our algorithm for definitional equality is based on joinability.

DEFINITION 3.1 (JOINABILITY). We say that two terms a and b are joinable (denoted as $a \downarrow b$) if there exists some term c such that $a \sim_{\beta n}^{*} c$ and $b \sim_{\beta n}^{*} c$.

A direct consequence of the $\beta\eta$ -confluence result is that $\beta\eta$ -joinability is transitive.

LEMMA 3.12 (TRANSITIVITY OF JOINABILITY). If $a, b, c \in SN$ and $a \downarrow b \downarrow c$, then $a \downarrow c$.

538 539

493

494

495 496

497

498

499 500

501

502

503

504 505

506

507

508 509

510

511

512

513 514

515

516

517

518

519 520

521

522

523

524

525 526

527

528

529

531

532

533

536

537

PROOF. Immediate by Theorem 3.1, Corollary 3.2, and Lemma 3.3.

Furthermore, unlike $\beta\eta$ -equivalence, the injectivity of neutral forms for joinability is immediate from its definition.

Lemma 3.13 (Injectivity of $\beta\eta$ -joinability).

- If $e_0 a_0 \downarrow e_1 a_1$, then $e_0 \downarrow e_1$ and $a_0 \downarrow a_1$.
- If $\pi_1 e_0 \downarrow \pi_1 e_1$ or $\pi_2 e_0 \downarrow \pi_2 e_1$, then $e_0 \downarrow e_1$.

Because joinability is reflexive and symmetric by definition, and transitive through the reasoning above, we can conclude that it is indeed an equivalence relation. Thus, we can use it as a syntactic model for declarative equality by simply erasing typing information.

Next, we define a reduce-and-compare algorithm for *subtyping*. First, consider the following untyped *one-step* subtyping relation, which compares its inputs without using $\beta\eta$ -reduction.

 $553 \qquad A \leq_1 B$

(Untyped reductionless subtyping)

0 D	S-Univ	S-Pi		S-Sig	
S-Refl	$i \leq j$	$A_1 \leq_1 A_0$	$B_0 \leq_1 B_1$	$A_0 \leq_1 A_1$	$B_0 \leq_1 B_1$
$\overline{A \leq_1 A}$	$\overline{\mathcal{U}_i \leq_1 \mathcal{U}_j}$	$\overline{\Pi x : A_0 . B_0} \leq$	$x_1 \Pi x : A_1 . B_1$	$\overline{\Sigma x : A_0. B_0} \leq$	$\leq_1 \Sigma x : A_1 . B_1$

⁵⁵⁸ We extend this definition to an algorithm for subtyping which first $\beta\eta$ -reduces its inputs.

DEFINITION 3.2 (UNTYPED SUBTYPING). We say that A is an untyped subtype of B, which we denote as $A \leq B$, if there exists some terms C_0 and C_1 such that $A \sim_{\beta_n}^* C_0$, $B \sim_{\beta_n}^* C_1$, and $C_0 \leq_1 C_1$.

We can also prove that untyped subtyping is transitive, after showing a few structural properties.

LEMMA 3.14 (SUB1 TRANSITIVE). If $A \leq_1 B$ and $B \leq_1 C$, then $A \leq_1 C$.

PROOF. We prove the following more general statement:

$$\forall A B, A \leq_1 B \implies \forall C, (B \leq_1 C \implies A \leq_1 C) \land (C \leq_1 A \implies C \leq_1 B)$$

The proof itself is straightforward by induction over the derivation of $A \leq_1 B$.

LEMMA 3.15 (SUB1 COMMUTATIVITY). If $A \leq_1 B$ and $A \sim_{\beta\eta} A_0$, then $\exists B_0, B \sim_{\beta\eta} B_0 \land A_0 \leq_1 B_0$.

LEMMA 3.16 (SUB1 PRESERVES SN). Let A, B, and C be arbitrary lambda terms. The following statements hold.

• $A \in \mathbf{SNe} \land (A \leq_1 B \lor B \leq_1 A) \implies A = B$

•
$$A \in SN \land (A \leq_1 B \lor B \leq_1 A) \implies B \in SN$$

• $A \rightsquigarrow_{SN} B \land (A \leq_1 B \lor B \leq_1 A) \implies A = B$

PROOF. By mutual induction over the SN definitions.

The transitivity of untyped subtyping then follows as a corollary of the confluence of $\beta\eta$ -reduction and the above lemmas.

COROLLARY 3.3 (TRANSITIVITY OF SUBTYPING). If A, B, $C \in SN$ and $A \leq B \leq C$, then $A \leq C$.

The benefit of using the directed definition of untyped subtyping ($A \le B$) over a declarative version obtained by erasing type annotations from $\Gamma \vdash A \le B$ is that the following no-confusion and injectivity properties are immediate.

LEMMA 3.17 (NOCONFUSION FOR UNTYPED SUBTYPING). Suppose $A, B \in SN$ and $A \leq B$, then the following holds:

- *A* and *B* cannot be type constructors with distinct head forms.
 - If A is in weak head neutral form, then B cannot be a type constructor.
 - If A and B are both weak head neutral forms, then they cannot have distinct elimination forms.

Lemma 3.18 (untyped injectivity of type constructors).

- If $\Pi x: A_0. B_0 \leq \Pi x: A_1. B_1$, then $A_1 \leq A_0$ and $B_0 \leq B_1$.
- If $\Sigma x: A_0$. $B_0 \le \Sigma x: A_1$. B_1 , then $A_0 \le A_1$ and $B_0 \le B_1$.
- If $\mathcal{U}_i \leq \mathcal{U}_j$, then $i \leq j$.

We use the transitivity, injectivity, and no-confusion properties of joinability and untyped subtyping in the next subsection to show that they are adequate models of the declarative relations.

3.7 Strong Normalization

 $[\![A]\!]_i \searrow S$

(Logical predicate)

$$\frac{SWT-PI}{F \in \mathbf{Term} \to \mathcal{P}(\mathbf{Term}) \quad [\![A]\!]_i \searrow S_A \quad \forall a, a \in S_A \to [\![B[a/x]]\!]_i \searrow F(a)}{[\![\Pi x : A. B]\!]_i \searrow \{b \mid b \in \hat{\Pi}(S_A, F)\}}$$

$$\frac{F \in \mathbf{Term} \to \mathcal{P}(\mathbf{Term}) \quad [\![A]\!]_i \searrow S_A \quad \forall a, a \in S_A \to [\![B[a/x]]\!]_i \searrow F(a)}{[\![\Sigma x : A. B]\!]_i \searrow \{c \mid c \in \hat{\Sigma}(S_A, F)\}}$$

$$\frac{A \sim_{SN} B \quad [\![B]\!]_i \searrow S}{[\![A]\!]_i \searrow S} \qquad \qquad \begin{array}{c} SWT\text{-}UNIV \\ j < i \\ \hline \hline \begin{bmatrix} \mathcal{U}_j \end{bmatrix} \\ i \searrow \{A \mid \exists S, [\![A]\!]_j \searrow S\} \end{array} \qquad \qquad \begin{array}{c} SWT\text{-}NE \\ A \in SNe \\ \hline \hline \begin{bmatrix} \mathcal{U}_j \end{bmatrix} \\ i \searrow \{A \mid \exists S, [\![A]\!]_j \searrow S\} \end{array} \qquad \qquad \begin{array}{c} SWT\text{-}NE \\ \hline \hline \begin{bmatrix} A \end{bmatrix} \\ i \searrow \{a \mid a \in SNe*\} \\ (Auxiliary definitions) \end{array}$$

$$a \in \mathbf{SNe}^* \triangleq \exists b, a \rightsquigarrow_{\mathbf{SN}}^* b \land b \in \mathbf{SNe}$$

$$b \in \widehat{\Pi}(S_A, F) \triangleq \forall a, a \in S_A \to b \ a \in F \ (a)$$

$$c \in \widehat{\Sigma}(S_A, F) \triangleq c \in \mathbf{SNe}^* \lor \exists a \ b, c \rightsquigarrow_{\mathbf{SN}}^* (a, b) \land a \in S_A \land b \in F(a)$$

Fig. 7. Logical Predicate

Finally, we use a logical predicate to show that all syntactically well-typed terms are contained in SN, i.e. are strongly normalizing.

Inspired by the inductive model for universes from Abel et al. [2008], we apply the *SN*-method [Joachimski and Matthes 2003] and define our logical predicate as an inductive relation over untyped lambda terms, closed by \sim_{SN} in rules SWT-SIGMA, SWT-STEP, and SWT-NE.

The logical predicate (Figure 7) takes the form $[\![A]\!]_i \searrow S$, meaning that the raw term *A* is a semantically well-formed type from the *i*th universe and is interpreted as the set of lambda terms *S*. Note that Rocq does not support the definition from Figure 7 as is because it requires nesting an inductive definition inside a recursive function over the universe level *i*. Our Rocq encoding is based on Liu et al. [2024a], which defines the predicate inductively over a parameterized interpretation of lower universes and then ties the knot by well-founded recursion over the universe level *i*.

An immediate result we can show is adequacy, which implies that every well-formed type is inhabited by terms from **SNe**, and every term that lives in the logical predicate is in the set *SN*.

LEMMA 3.19 (ADEQUACY). If $[A]_i \searrow S$, then SNe $\subseteq S \subseteq$ SN and $A \in$ SN.

PROOF. By nested induction over the universe level *i* and the derivation of
$$[A]_i \searrow S$$
. The only
interesting cases are SWT-P1 and SWT-SIGMA, both of which follow from Lemma 3.5.
The interpreted set *S* is always closed under $a \rightarrow_{SN} b$.
LEMMA 3.20 (BACKWARD CLOSURE). *If* $[A]_i \searrow S, b \in S$, and $a \rightarrow_{SN} b$, then $a \in S$.
LEMMA 3.21 (LOGICAL PREDICATE CASES). *If* $[[A]_i \searrow S$, then there exists some type constructor A_0
such that $A \sim_{SN}^{*} A_0$ and $[A]_i \searrow S$.
PROOF. By straightforward induction over the derivation of $[A]_i \searrow S$.
The following lemma semantically justifies the untyped subtyping relation.
LEMMA 3.22 (LOGICAL PREDICATE PRESERVED BY SUBTYPING).
 $\forall A B S_0 S_1, [A]_i \searrow S_0 \land [B]_i \searrow S_1 \land A \leq B \implies S_0 \subseteq S_1$
PROOF. As in Lemma 3.14, we need to prove the following generalized statement.
 $\forall A B S_0 S_1, [A]_i \searrow S_0 \land [B]_i \searrow S_1 \implies (A \leq B \implies S_0 \subseteq S_1) \land (B \leq A \implies S_1 \subseteq S_0)$
The proof proceeds by induction over the derivation of $[A]_i \searrow S_0$ followed by case analyzing
the derivation of $[B]_i \searrow S$ with the help of Lemma 3.21. Lemma 3.17 is used to rule out the
contradictory cases where *A* and *B* evaluates to distinct head forms. Lemma 3.18 is used to obtain
the subtyping premises required to apply the induction hypotheses.

The next corollary allows us to treat our logical predicate as a partial function that takes a type
and universe as inputs, and returns a set of lambda terms as its output.

COROLLARY 3.4 (LOGICAL PREDICATE IS FUNCTIONAL).
 $\forall A S_0 S_1, [A]_i \searrow S_0 \land [A]_i \searrow S_1 \implies S_0 = S_1$
PROOF. Immediate by instantiating Lemma 3.22 with $A \leq A$.
We can prove that the logical predicate is cumulative.

LEMMA 3.23 (LOGICAL PREDICATE CUMULATIVITY). *If* $[A]_i \searrow S$ and $i \leq j$, then $[A]_j \searrow S$.

PROOF. Immediate by induction over the derivation of $[A]_i \searrow S$.

PROOF. Immediate by induction over the derivation of $[A]_i \searrow S$.

PROOF. Immediate by induction over the derivation of $[A]_i \searrow S$.

PROOF. Immediate by induction over the derivation of $[A]_i \gtrsim S$.

PROOF. Immediate by ind

⁶⁸⁷ DEFINITION 3.5 (SEMANTIC EQUALITY). The terms a and b are semantically equal with type A ⁶⁸⁸ under the context Γ , denoted as $\Gamma \models a = b : A$, when $a \downarrow b$ and $\Gamma \models a : A$ and $\Gamma \models b : A$.

DEFINITION 3.6 (SEMANTIC SUBTYPING). The term A is a semantic subtype of B under the context Γ , denoted as $\Gamma \models A \leq B$, when $A \leq B$ and $\Gamma \models a : A$ and $\Gamma \models b : A$.

DEFINITION 3.7 (SEMANTIC CONTEXT WELL-FORMEDNESS). The context Γ is semantically well-formed, denoted as $\models \Gamma$, when for all $x : A \in \Gamma$, there exists some universe level i such that $\Gamma \models A : \mathcal{U}_i$.

By unfolding the definitions above, we immediately obtain the following structural lemmas about semantic typing.

697 LEMMA 3.24 (SEMANTIC WEAKENING). If $\Gamma \models a : A$ and $\Gamma \models B : \mathcal{U}_i$ then $\Gamma, x : B \models a : A$

Lemma 3.25 (Semantic Substitution). If $\Gamma \models a : A$ and $\Gamma, x : A \models b : B$ then $\Gamma \models b[a/x] : B[a/x]$

Similar results for equality and subtyping are omitted as they directly follow from the above.
 Context well-formedness satisfies the following structural lemmas.

Lemma 3.26 (Structural rules for semantic well-formedness).

 $\models \Gamma \text{ and } \Gamma \models A : \mathcal{U}_i \text{ implies } \models \Gamma, x : A$

PROOF. The empty context case is trivial. The cons case requires Lemma 3.24.

Note that all of the structural rules we have seen so far can be proven by unfolding the definition of semantic typing. Thus, they are useful for organizing the proof but are not necessary for deriving the fundamental theorem.

THEOREM 3.2 (FUNDAMENTAL THEOREM).

⊧・

- $\vdash \Gamma$ implies $\models \Gamma$
- $\Gamma \vdash a : A \text{ implies } \Gamma \models a : A$
- $\Gamma \vdash a = b : A \text{ implies } \Gamma \models a = b : A$
- $\Gamma \vdash A \leq B$ implies $\Gamma \models A \leq B$

PROOF. By mutual induction over the typing judgments and the structural lemmas.

COROLLARY 3.5 (Well-TYPED STRONG NORMALIZATION). If $\Gamma \vdash a : A$ then $a \in SN$ and $A \in SN$. COROLLARY 3.6 (COMPLETENESS OF REDUCE-AND-COMPARE).

- $\Gamma \vdash A \leq B$ implies $A \leq B$
- $\Gamma \vdash a = b : A \text{ implies } a \downarrow b$

4 Total Correctness via Coquand's Algorithm

Corollary 3.6 gives us the completeness of the reduce-and-compare algorithm with respect to typed definitional equality and subtyping. To prove its soundness, a sufficient condition is to show that $\beta\eta$ -reduction is type-preserving. Unfortunately, $\beta\eta$ -reduction is known to violate subject reduction [Abel and Coquand 2007; Lennon-Bertrand 2022].

However, η -expansion, based on either the shape of the term or its type, is type-preserving. As a result, it is easier to show the soundness of algorithms that perform η -expansion with respect to typed convertibility. Therefore, we next use Coquand's algorithm extended with surjective pairing [Abel and Coquand 2007; Coquand 1991], which is based on η -expansion, as a bridge to show the soundness of our normalize-and-compare algorithm.

This detour through Coquand's algorithm is itself worthwhile as we gain a correctness proof for
 a second decision procedure, which can be more efficient.

735

689

690

691

692

693 694

695

696

702

703

704

705

706

707

708

709 710

711

712

713

714

715

720

721 722

723

П

Definitional Subtyping ($\Gamma \vdash A \leq B$) Corollary 3.6 Reduce-and-Compare ($\Gamma \vdash A, B : \mathcal{U}_i \land A \leq B$) Lemma 4.8 Coquand's Algorithm ($\Gamma \vdash A, B : \mathcal{U}_i \land A \ll B$)

Fig. 8. Equivalence of Algorithmic and Declarative Subtyping Relations

In this section, we first extend the syntactic metatheory of Coquand's algorithm [Goguen 2005] with pairs and subtyping. We then show that this extended version of Coquand's algorithm is sound with respect to typed convertibility (Lemma 4.4), complete with respect to untyped subtyping (Lemma 4.8), and terminates on SN terms (Lemma 4.6). The soundness and completeness results allows us to show that Coquand's algorithm, the reduce-and-compare algorithm, and typed convertibility are equivalent on well-typed terms, following the diagram in Figure 8. This equivalence result, combined with termination, completes our main decidability proof.

The most involved part of this process is the completeness proof (Lemma 4.8). We need to convert a potentially non-type preserving $\beta\eta$ -reduction sequence into a run of Coquand's algorithm: β reductions interleaved with η -expansions. Goguen [2005] shows that Coquand's algorithm (with only function η -law) can be syntactically justified from the confluence of $\beta\eta$ -reduction. We apply this technique to justify the completeness of both the subtyping and equality relation.

4.1 Coquand's Algorithmic Conversion by η -Expansion

Our version of Coquand's algorithmic equality, based on Coquand [1991] and Abel and Coquand [2007], is as follows.

(Algorithmic equality)

$$\frac{\substack{a \sim _{h}^{*} f_{0} \qquad b \sim _{h}^{*} f_{1} \qquad f_{0} \sim f_{1}}{a \leftrightarrow b}}{a \leftrightarrow b}$$

(Algorithmic equality for head normal forms)

	$r a \sim \lambda r h$	dr	а ~ P		dr h
CE-VAR CI	$\begin{array}{c} \text{E-Abs} \\ a \leftrightarrow b \end{array}$	$\begin{array}{l} \text{CE-AbsNeu} \\ a \leftrightarrow e x \end{array}$	$x \notin \mathbf{freevar}(a)$	$\begin{array}{l} \text{CE-NeuAbs} \\ e x \leftrightarrow b \end{array}$	$x \notin \mathbf{freevar}(a)$

CE-PAIR		CE-PAIRNEU		CE-NeuPair	CE-UNIX	
$a_0 \leftrightarrow a_1$	$b_0 \leftrightarrow b_1$	$a_0 \leftrightarrow \pi_1 e$	$b_0 \leftrightarrow \pi_2 e$	$\pi_1 \ e \leftrightarrow a_0$	$\pi_2 \ e \leftrightarrow b_0$	CL-UNIV
$(a_0, b_0) \sim$	(a_1, b_1)	(a_0, b_0)) ~ e	$e \sim (a)$	(a_0, b_0)	$\overline{\mathcal{U}_i \sim \mathcal{U}_i}$
CE-PI		CE-App	0	CE-Fst	CE-	Snd
$A_0 \leftrightarrow A$	$A_1 \qquad B_0 \leftrightarrow B_1$	$e_0 \sim e_1$	$a_0 \leftrightarrow a_1$	$e_0 \sim e$	² 1	$e_0 \sim e_1$
$\Pi x: A_0.$	$B_0 \sim \Pi x : A_1 . B_1$	e_0	$a_0 \sim e_1 a_1$	$\overline{\pi_1} e_0 \sim \pi$	$\overline{t_1 e_1} \qquad \overline{\pi_2}$	$e_0 \sim \pi_2 e_1$

We follow Goguen [2005] and split algorithmic equality into two mutually defined relations, where $a \leftrightarrow b$ is the top-level definition for algorithmic equality with a single rule that evaluates a and binto their respective weak head normal forms f_0 and f_1 , and then check that $f_0 \sim f_1$ through the auxiliary relation defined only for weak head normal forms.

 $a \leftrightarrow b$

 $f_0 \sim f_1$

Inspired by Coquand's equality algorithm, we define the algorithmic subtyping relation below, which performs weak-head reduction and compares types in weak-head normal forms.

787 788

 $A \ll B$

 $f_0 \leq f_1$

785

786

789 790 791

792

793

798

799

800

801 802

803

808

809

814

815 816

817

818

819 820

821

822 823

824

825

826

827

828

829

831

$$\frac{A \sim_{h}^{*} f_{0}}{A \ll_{h}^{*} f_{0}} \qquad B \sim_{h}^{*} f_{1} \qquad f_{0} \lesssim f_{1}}{A \ll B}$$

(Algorithmic subtyping for head normal forms)

 $\begin{array}{c} (\operatorname{FigorMinine} \operatorname{subtyping}) \text{ or head norm} \\ \hline CLE-UNIV \\ \hline u_i \leq \mathcal{U}_j \end{array} \qquad \begin{array}{c} \operatorname{CLE-PI} \\ A_1 \ll A_0 \\ \overline{\mathcal{U}_i \leq \mathcal{U}_j} \end{array} \qquad \begin{array}{c} \operatorname{CLE-Sig} \\ A_1 \ll A_0 \\ \overline{\mathcal{B}_0 \ll B_1} \\ \overline{\mathcal{I}_{X}:A_0,B_0 \leq \Pi x:A_1,B_1} \end{array} \qquad \begin{array}{c} \operatorname{CLE-Sig} \\ A_0 \ll A_1 \\ \overline{\mathcal{D}_{X}:A_0,B_0 \leq \Sigma x:A_1,B_1} \end{array} \qquad \begin{array}{c} \operatorname{CLE-Ner} \\ e_0 \sim e_1 \\ \overline{\mathcal{L}_{e_0} \leq e_1} \end{array}$ CLE-NEUNEU

Above, $A \ll B$ is the top-level subtyping relation defined over arbitrary terms and $f_0 \leq f_1$ the auxiliary subtyping relation over weak head normal forms. For convenience, we refer to both the equality and subtyping relations as Coquand's algorithm.

Coquand's algorithm uses the shape of the two terms being compared to guide η -expansion. In rule CE-AbsNeu, we have a lambda term λx . *a* on the left-hand side and a weak neutral term *e* on the right-hand side. Thus, the only way to make progress is to η -expand e into λx . e x so we can check the equality by comparing its body e x to a, the body of λx . a.

The subtyping relation assumes that the terms being related are types. In cases where either side 804 of $A \ll B$ reduce to a term constructor such as a lambda term or a pair, the algorithm returns false. 805 When both A and B are in weak-head neutral form, CLE-NEUNEU falls back to checking whether A 806 807 and *B* are equal neutral terms.

4.2 Soundness of Coquand's Algorithm

The soundness property states that well-typed terms related by algorithmic conversion must also 810 be convertible by the type-annotated subtyping relation. The proof of soundness, at a high-level, 811 812 amounts to showing that given $\Gamma \vdash A : \mathcal{U}_i, \Gamma \vdash B : \mathcal{U}_i$ and $A \ll B$, only well-typed terms appear in the derivation tree of $A \ll B$. 813

Before we dive into the soundness proof, we state two corollaries of the normalization property (Corollary 3.5) that we use to reason about subtyping.

LEMMA 4.1 (PI SUBTYPING). If $\Gamma \vdash \Pi x : A, B \leq C$ or $\Gamma \vdash C \leq \Pi x : A, B$, then there exists some A_0, B_0 , and i such that $\Gamma \vdash C = \Pi x : A_0 . B_0 : \mathcal{U}_i$.

LEMMA 4.2 (UNIV SUBTYPING). If $\Gamma \vdash \mathcal{U}_i \leq C$ or $\Gamma \vdash C \leq \mathcal{U}_i$, then there exists some *j* and *k* such that $\Gamma \vdash C = \mathcal{U}_i : \mathcal{U}_k$.

Now we can show the soundness of Coquand's equality with respect to definitional equality.

LEMMA 4.3 (SOUNDNESS FOR ALGORITHMIC EQUALITY).

• If $a \leftrightarrow b$ and $\Gamma \vdash a : A$ and $\Gamma \vdash b : A$, then $\Gamma \vdash a = b : A$.

• If $f_0 \sim f_1$, then

- $if \Gamma \vdash f_0 : A and \Gamma \vdash f_1 : A, then \Gamma \vdash f_0 = f_1 : A.$
- if $f_0 \sim f_1$ and both f_0 and f_1 are neutral, then if $\Gamma \vdash f_0 : A$ and $\Gamma \vdash f_1 : B$, there exists some *C* such that $\Gamma \vdash C \leq A$, $\Gamma \vdash C \leq B$, and $\Gamma \vdash f_0 = f_1 : C$.

PROOF. By mutual induction over the derivations of $a \leftrightarrow b$ and $f_0 \sim f_1$. We consider only case 830 CE-App to highlight complexities introduced by subtyping. Suppose we have $f_0 = e_0 a_0 \sim e_1 a_1 = f_1$ It suffices to show that the third bullet point holds as it subsumes the second bullet point when 832

833

(Algorithmic subtyping)

both terms are neutral. Suppose also that $\Gamma \vdash e_0 a_0 : A$ and $\Gamma \vdash e_1 a_1 : B$ for some A and B. By the generation lemma (Lemma 2.2), there exists some terms A_0, B_0, A_1 and B_1 such that $\Gamma \vdash e_0 : \Pi x :$ $A_0, B_0, \Gamma \vdash e_1 : \Pi x : A_1, B_1, \Gamma \vdash a_0 : A_0, \Gamma \vdash a_1 : A_1, \Gamma \vdash B_0[a_0/x] \le A$ and $\Gamma \vdash B_1[a_1/x] \le B$.

Applying the induction hypothesis to the neutral forms e_0 and e_1 , we know that there exists some type *C* such that $\Gamma \vdash e_0 = e_1 : C$ where $\Gamma \vdash C \leq \Pi x : A_0 . B_0$ and $\Gamma \vdash C \leq \Pi x : A_1 . B_1$.

By Lemma 4.1, there exists types A_2 , B_2 and i such that $\Gamma \vdash \Pi x : A_2$. $B_2 = C : \mathcal{U}_i$.

Therefore, we can replace the occurrences of *C* with $\Pi x: A_2. B_2$, obtaining the following statements.

- $\Gamma \vdash e_0 = e_1 : \Pi x : A_2 . B_2$
- $\Gamma \vdash \Pi x : A_2 . B_2 \leq \Pi x : A_0 . B_0$
- $\Gamma \vdash \Pi x : A_2 . B_2 \leq \Pi x : A_1 . B_1$

By the injectivity rule L-PIPROJONE, we deduce $\Gamma \vdash A_0 \leq A_2$ and $\Gamma \vdash A_1 \leq A_2$. This allows us to convert the type of both a_0 and a_1 to A_2 , so we can obtain $\Gamma \vdash a_0 = a_1 : A_2$ from the induction hypothesis. By rule E-APP, we have $\Gamma \vdash e_0 a_0 = e_1 a_1 : B_2[a_0/x]$. It now suffices to show that $\Gamma \vdash B_2[a_0/x] \leq B_0[a_0/x]$ and $\Gamma \vdash B_2[a_0/x] \leq B_1[a_1/x]$.

We can prove $\Gamma \vdash B_2[a_0/x] \leq B_0[a_0/x]$ immediately by L-PIPRoJTwo. To prove $\Gamma \vdash B_2[a_0/x] \leq B_1[a_1/x]$, we cannot directly apply rule L-PIPRoJTwo as a_0 is not necessarily typed under A_1 . Instead, we need to first apply transitivity and show $\Gamma \vdash B_2[a_0/x] \leq B_2[a_1/x] \leq B_1[a_1/x]$. The two subgoals then follow from L-PIPRoJTwo.

The soundness property for subtyping is formulated similarly, though we do not need to include the special case when both types *A* and *B* are neutral since subtyping for neutral types degenerate into equality, whose soundness we have already proven.

Lemma 4.4 (Soundness for Algorithmic Subtyping).

- If $A \ll B$, and $\Gamma \vdash A : \mathcal{U}_i$ and $\Gamma \vdash B : \mathcal{U}_i$, then $\Gamma \vdash A \leq B$.
- If $f_0 \leq f_1$, and $\Gamma \vdash f_0 : \mathcal{U}_i$ and $\Gamma \vdash f_1 : \mathcal{U}_i$, then $\Gamma \vdash f_0 \leq f_1$.

PROOF. By mutual induction on $A \leq B$ and $A \ll B$. The only interesting case is CLE-NEUNEU, where *A* and *B* are neutral. We finish the proof by simply composing Lemma 4.3 and rule L-Eq. \Box

4.3 Termination of Coquand's Algorithm

To prove termination, we use a termination metric over the input terms *a* and *b* from Goguen [2005], extended to include pairs. (A similar metric with pairs but only for β -forms can be found in Abel and Coquand [2007].)

DEFINITION 4.1 (TERMINATION METRIC FOR COQUAND'S ALGORITHM). Given two β -normalizing terms a and b, there must exists some β -normal forms u_0 and u_1 such that $a \sim_{l_0}^* u_0$ and $b \sim_{l_0}^* u_1$. Let m and n be the number of steps it takes for a and b to reduce to their respective normal form, we define the termination metric $\mathcal{T}(a, b) = m + n + |u_0| + |u_1|$, where the $|\cdot|$ is a size function over lambda terms, defined as follows.

 $\begin{aligned} |\lambda x. a| &= 3 + |a| \\ |a b| &= 1 + |a| + |b| \\ |(a, b)| &= 3 + |a| + |b| \\ |\pi_1 a| &= 1 + |a| \\ |\pi_2 a| &= 1 + |a| \end{aligned}$

The $|\cdot|$ function weighs constructors more than elimination forms, allowing us to show that the inputs to recursive calls in CE-AbsNeu and CE-PAIRNeu are decreasing.

881 882

839

842

843

844

850

851

852

853

854

855

856 857

858

859

860

861

862 863

864

865

866

867 868

869

870

871

872

873

874

875

876

877

878 879

Goguen [2005] uses the definition of \mathcal{T} as metric for both the termination and completeness 883 proof. We take a more modular approach by introducing an intermediate inductive relation that 884 characterizes the domain of the conversion algorithm. The intermediate relation serves two purposes. 885 First, once we have shown that every pair of terms *a* and *b* such that $\mathcal{T}(a, b)$ is defined inhabit 886 the domain, then we can immediately recover termination by the Bove-Capretta method [Bove 887 and Capretta 2005]. Second, the intermediate relation abstracts away the low-level details of the 888 termination metric and allows us to reason by induction over the call-graph of the algorithm, 889 significantly simplifying our reasoning. 890

We write the relations encoding the domains for equality and subtyping as $(f_0, f_1) \in \mathcal{A}$ and $(f_0, f_1) \in \mathcal{S}$, which correspond to $f_0 \sim f_1$ and $f_0 \leq f_1$, and $(a, b) \in \mathcal{A}^*$ and $(A, B) \in \mathcal{S}^*$, which correspond to $a \leftrightarrow b$ and $A \ll B$. We omit the definition of these relations here as they are a straightforward application of the Bove-Capretta method. The definition can be found in the supplementary material.

The following properties can be proven by induction over the termination metric. We omit the proof, as it follows the exact same structure as the termination proof by Goguen [2005].

LEMMA 4.5 (METRIC IMPLIES DOMAIN). If there exists some n such that $\mathcal{T}(a, b) = n$, then $(a, b) \in \mathcal{A}^*$ and $(a, b) \in \mathcal{S}^*$.

By Corollary 3.5, all well-typed terms are strongly normalizing, and therefore for every pair of terms *a* and *b*, $\mathcal{T}(a, b)$ is defined. By Lemma 4.5, we deduce that $(a, b) \in \mathcal{A}^*$ and $(a, b) \in \mathcal{S}^*$, and we can then define the equality and subtyping functions recursively over the derivations of the domains, concluding the proof of the following termination result.

LEMMA 4.6 (TERMINATION OF COQUAND'S ALGORITHM). If $\Gamma \vdash A : \mathcal{U}_i$ and $\Gamma \vdash B : \mathcal{U}_i$, then $A \ll B$ terminates.

Furthermore, fixing *a* and *b* as inputs, we can show that the derivation of $(a, b) \in \mathcal{A}^*$ or $(a, b) \in \mathcal{S}^*$ does not affect the result of the algorithm. In our Rocq mechanization, we leverage this irrelevance property to place the domain definitions in the Prop sort so the extracted conversion algorithm does not have to compute the proofs or recurse over a gas parameter.

4.4 Completeness of Coquand's Algorithm w.r.t Untyped Subtyping

To show completeness, we reuse the domain definitions from Section 4.3 as our induction metric and formulate the completeness theorem as follows.

Lemma 4.7 (Completeness of Algorithmic Equality (auxiliary)).

- If $(a, b) \in \mathcal{A}^*$, $\Gamma \vdash a : A$, $\Gamma \vdash b : A$ and $a \downarrow b$, then $a \leftrightarrow b$. • If $(f, f) \in \mathcal{A}$ then
 - If $(f_0, f_1) \in \mathcal{A}$, then
 - *if* $\Gamma \vdash f_0 : A$, $\Gamma \vdash f_1 : A$ and $f_0 \downarrow f_1$, then $f_0 \sim f_1$.
 - if f_0 and f_1 are both weak head neutral forms, then given $\Gamma \vdash f_0 : A, \Gamma \vdash f_1 : B$ and $f_0 \downarrow f_1$, we have $f_0 \sim f_1$.

PROOF. By mutual induction over the derivations of $(a, b) \in \mathcal{A}^*$ and $(f_0, f_1) \in \mathcal{A}$. The proof obligations can be split into two categories: proving the well-typedness of the subterms to be able to invoke the inductive hypotheses, and showing that the recursive calls are made on terms are still $\beta\eta$ -joinable so we can meet the prerequisite for applying the induction hypotheses.

We consider the case for rule A-APPCONG as an example. In case A-APPCONG, we have $(e_0 \ a_0, e_1 \ a_1) \in \mathcal{A}$ where $(a_0, a_1) \in \mathcal{A}^*$, $(e_0, e_1) \in \mathcal{A}$, $\Gamma \vdash e_0 \ a_0 : C_0$, $\Gamma \vdash e_1 \ a_1 : C_1$, and $e_0 \ a_0 \downarrow e_1 \ a_1$. By the Generation Lemma (Lemma 2.2), there exists some $A_0 \ A_1 \ B_0 \ B_1$ such that $\Gamma \vdash e_0 : \Pi x : A_0 . B_0$, $\Gamma \vdash a_0 : A_0, \ \Gamma \vdash e_1 : \Pi x : A_1 . B_1, \ \Gamma \vdash a_1 : A_1$.

896

897 898

899

900

901

902

903

904 905

906

907 908

909

910

911

912

913

914

915

916

917

918

919

920

921

922

923

924

925

By the injectivity of application forms for $\beta\eta$ -joinability (Lemma 3.13), we deduce from $e_0 a_0 \downarrow e_1 a_1$ that $e_0 \downarrow e_1$ and $a_0 \downarrow a_1$. Since we also know that $\Gamma \vdash e_0 : \Pi x : A_1 . B_1, \Gamma \vdash e_1 : \Pi x : A_0 . B_0$ from the generation lemma, we can invoke the induction hypothesis and derive $e_0 \sim e_1$.

By Lemma 4.3, we know that there exists a common subtype *C* for $\Pi x: A_0$. B_0 and $\Pi x: A_1$. B_1 . By Lemma 4.1 and rule L-PIPROJONE, there exists some A_2 such that $\Gamma \vdash A_0 \leq A_2$, $\Gamma \vdash A_1 \leq A_2$. Thus, we have $\Gamma \vdash a_0 : A_2$, $\Gamma \vdash a_1 : A_2$, and $a_0 \downarrow a_1$. By the induction hypothesis, we deduce $a_0 \leftrightarrow a_1$. From $e_0 \sim e_1$ and $a_0 \leftrightarrow a_1$, we conclude that $e_0 a_0 \sim e_1 a_1$.

The completeness of algorithmic subtyping is stated as follows.

LEMMA 4.8 (COMPLETENESS OF COQUAND'S ALGORITHMIC SUBTYPING (AUXILIARY)).

• If $(f_0, f_1) \in S$, $\Gamma \vdash f_0 : \mathcal{U}_i$, $\Gamma \vdash f_1 : \mathcal{U}_i$, and $f_0 \leq f_1$, then $f_0 \leq f_1$.

• If $(A, B) \in S^*$, $\Gamma \vdash A : \mathcal{U}_i, \Gamma \vdash B : \mathcal{U}_i$, and $A \leq B$, then $A \ll B$.

PROOF. By mutual induction over the derivations of $(A, B) \in S^*$ and $(A, B) \in S$. The neutral case (CLE-NEUNEU) requires Lemma 4.7.

Composing the completeness of Coquand's algorithm with respect to untyped subtyping (Lemma 4.8), the soundness of Coquand's algorithm with respect to definitional subtyping (Lemma 4.4), and the completeness of untyped subtyping with respect to definitional subtyping (Corollary 3.6), we conclude our decidability proof with the following equivalence results.

LEMMA 4.9 (REDUCE-AND-COMPARE AND COQUAND'S ALGORITHM ARE CORRECT). Given $\Gamma \vdash A : \mathcal{U}_i$ and $\Gamma \vdash B : \mathcal{U}_j$, the statements $\Gamma \vdash A \leq B$, $A \leq B$, and $A \ll B$ are all equivalent.

THEOREM 4.1 (TYPE CONVERSION IS DECIDABLE). Given $\Gamma \vdash A : \mathcal{U}_i$ and $\Gamma \vdash B : \mathcal{U}_j$, the relation $\Gamma \vdash A \leq B$ is decidable.

5 Discussion

5.1 Admissibility of **Π**-Injectivity

One nonstandard feature of $\lambda^{\Pi, \Sigma, U_i, \mathbb{N}}$ is the inclusion of the injectivity rules L-PIPROJ1 and L-PIPROJ2. We explicitly add these rules as part of the definitional subtyping relation so that we can prove subject reduction syntactically, simplifying the definition of our logical predicate in Section 3.7.

These injectivity rules are sound and justified by our algorithm so they do not compromise the implementation of type checking. Indeed, the results from Section 4 allow us to replace conversion rule WT-CONV with the following rule based on algorithmic subtyping.

$$\frac{\Gamma \vdash a : A}{\Gamma \vdash a : B} \frac{\Gamma \vdash B : \mathcal{U}_i \qquad A \ll B}{\Gamma \vdash a : B}$$

Algorithmic subtyping does not depend on definitional equality, so the inclusion of injectivity rules in the latter does not affect the implementation of the type checker.

That said, it is difficult to remove the injectivity rules from the specification of the type system. While the rules do not affect the operation of algorithm, they are needed to show its correctness. The justification of rule WT-ACONV ultimately relies on subject reduction, which in turn depends on II-injectivity. Therefore, the correctness of rule WT-ACONV does not imply that II-injectivity is admissible.

Why would we want to show that Π-injectivity is admissible? After all, they causes no issues
with our implementation and we have shown that the type system with these rules is decidable.
The real issue lies in model construction. While the reducibility model in Section 4 satisfies the
injectivity of type constructors, there also exist models of dependent type theory where injectivity

of type constructors is not true. For example, the standard set theoretic models by Miquel and Werner [2003] and Timany and Sozeau [2017] forget too much information about the syntax of the type theory for injectivity to hold semantically. But, the contravariant subtyping rule for function types already rules out set theoretic models for $\lambda^{\Pi,\Sigma,U_i,\mathbb{N}}$. Therefore, we see no little reason to avoid the inclusion of the injectivity rules that simplify our proofs.

Alternatively, if the admissibility of type injectivity is needed, it is possible to replace the 986 logical predicate in Section 3.7 with a Kripke-style logical relation, from which injectivity of type 987 constructors and strong normalization can both be derived at once. The Kripke-style logical relation 988 can remain oblivious of the algorithmic conversion rules as the correctness of the algorithm can 989 be obtained by composing Π -injectivity and strong β -normalization with our syntactic proof in 990 Section 4. Thus, while we would still need to pay the cost of constructing a PER model to obtain 991 Π -injectivity, we can avoid the intricate setup of a parameterized generic equality found in Abel 992 et al. [2017]; Adjedj et al. [2024] and keep our Kripke-style logical relation minimal. 993

5.2 Annotated Lambda Terms

Another slightly nonstandard feature of $\lambda^{\prod, \Sigma, U_i, \mathbb{N}}$ is that function abstractions are not annotated with the types of their arguments. Such annotations are known to invalidate the confluence property for $\beta\eta$ -reduction. Suppose x, y, and z are distinct variables, the term $\lambda x : A.(\lambda y : B.z)x$ can either η -reduced to $\lambda y : B.z$ or β -reduced to $\lambda x : A.z$. Thus, when A and B are distinct types, we obtain a counterexample to confluence.

However, the type system is designed such that annotating the lambda terms with types should not lead to significant changes to our proof development. Suppose the terms $\lambda x : A.a$ and $\lambda x : B.b$ share the same type. By the generation lemma, we can find some common subtype *C* of the types *A* and *B*. If we continue to use rule E-ABSEXT to subsume both the congruence and η -law for abstractions, then there is no need to explicitly compare the annotations *A* and *B* as the congruence rule implied by rule E-ABSEXT identifies lambda terms whose domains share a common subtype. Thus, the $\beta\eta$ -confluence result over erased lambda terms is sufficient to justify the correctness of the algorithm, even though the syntax involves type-annotated lambdas. In our work, we consider only unannotated lambda terms so we do not have to prove more simulation/commutativity lemmas between erasure, one-step subtyping, and reduction.

What happens if the declarative system explicitly requires the annotations to be equal in the congruence rule for abstractions, in a way that is not supported by the generation lemma? For example, suppose we are working with a pure-type system without uniqueness of sorts, and we replace rule E-ABSEXT with the following congruence rule for type-annotated abstractions (cf. Siles and Herbelin [2012]).

$$\frac{\Gamma - ABSCONG}{\Gamma \vdash A = A_0 : \mathcal{U}_i \qquad \Gamma, x : A \vdash B : \mathcal{U}_i \qquad \Gamma, x : A \vdash a = b : B}{\Gamma \vdash \lambda x : A \cdot a = \lambda x : A_0 \cdot b : \Pi x : A \cdot B}$$

In that case, we could not use the $\beta\eta$ -confluence result over erased lambda terms. The conversion algorithm would be forced to compare annotations, but $\beta\eta$ -joinability over erased terms does not include enough information about the annotations.

5.3 Relaxing the SN Condition

Our confluence property for $\beta\eta$ -contraction holds only for the terms in the set SN. However, this condition is stronger than necessary.

994

995 996

997

998

999

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

1015

1022

1023

1024 1025

1026

We use SN in two ways. To derive η -postponement, we use the SN predicate to rule out terms that 1030 can step into stuck forms. To derive $\beta\eta$ -confluence, we also use the SN predicate to find β -normal 1031 forms. 1032

However, to rule out terms that can get stuck, it suffices to find a predicate P that satisfies 1033 following properties, which do not imply that terms are strongly normalizing: 1034

1035 1036 1037

1041

1042

1043

1044

1045 1046

1047

1048

1049 1050

- If *P*(*a*) holds, then *a* is not a stuck term.
- If *P*(*a*) holds, then *P*(*b*) holds if *b* is a subterm of *a*.
 - If P(a) holds and $a \rightsquigarrow_{\beta\eta} b$, then P(b) holds.

1038 In our Rocq development, our η -postponement property is parameterized over such a predicate P. 1039 An alternative to SN is a predicate that holds precisely when terms cannot reduce to a stuck 1040

state. Consider the following coinductively defined set.

DEFINITION 5.1. We define okay as the the largest set of untyped lambda terms that satisfies the following properties.

- If $a \in okay$, then a does not contain any stuck subterms.
- If $a \in \mathbf{okay}$ and $a \rightsquigarrow_{\beta\eta} b$, then $b \in \mathbf{okay}$.

Note that **okay** contains the non-terminating term $(\lambda x. x x) (\lambda x. x x)$ since it loops without ever reaching a stuck state.

PROPOSITION 5.1. Terms in the set okay satisfy the postponement of η -reductions.

PROOF. It suffices to show that **okay** satisfies the three condition we specify above.

1051 The first and third properties are immediate by definition. The second property can be verified 1052 by the coinduction principle. П 1053

Thus, the η -postponement proof can be extended to systems that are not necessarily terminating as long as the terms do not get stuck. In particular, showing that all well-typed terms are okay, a similar property to type safety, is sufficient.

To derive $\beta\eta$ -confluence from η -postponement, we do not need all reduction sequences to 1057 result in a β -normal form, we only need *one* reduction sequence to terminate. Thus, assuming 1058 η -postponement, we can obtain $\beta\eta$ -confluence from weak β -normalization. 1059

6 **Related Work**

6.1 Syntactic Methods for Type Conversion 1062

In this section, we compare our approach with other syntactic proofs of decidable type conversion for 1063 typed $\beta\eta$ -equality. (In systems with *untyped* β -equivalence as definitional equality, the decidability 1064 of type conversion follows immediately from strong normalization and confluence [Barendregt 1065 1993; Geuvers 1994; Luo 1990].) 1066

Goguen [2005] proves the decidability of a logical framework with a typed $\beta\eta$ -equality as its 1067 convertibility relation. This system contains only the η -law for functions; surjective pairing is not 1068 considered. Unlike $\lambda^{\Pi,\Sigma,U_i,\mathbb{N}}$, the logical framework does not support large eliminations or polymor-1069 phism. Goguen shows that the soundness, completeness, and termination of Coquand's untyped 1070 algorithm [Coquand 1991] and Pfenning and Harper's typed algorithm [Harper and Pfenning 1071 2005] can be justified by a syntactic proof parameterized over the confluence of $\beta\eta$ -reduction and 1072 metatheoretic properties about the type system, including normalization and subject reduction (for 1073 $\beta\eta$ -reduction). Instead of proving completeness with respect to the typed definitional equality, he 1074 proves the completeness of algorithmic conversion with respect to untyped $\beta\eta$ -equivalence. That 1075 is, if two well-typed terms $\beta\eta$ -reduce to the same term, then one can show their equivalence by 1076 performing η -expansion and β -reduction. 1077

1054

1055

1056

1060

1061

1078

In our work, we make the observation that we can perform Goguen's completeness proof without requiring η -reduction to be type-preserving. This allows us to use the completeness proof as a bridge to relate untyped $\beta\eta$ -reduction, which can relate ill-typed terms, to the typed convertibility relation, which only relates well-typed terms. Furthermore, our development applies to a more expressive language, which includes large eliminations, surjective pairing, and subtyping. The latter two features mean that we must rely on SN to show confluence and that our development cannot rely on the uniqueness of typing.

1086 Abel and Coquand [2007] use a PER model to justify two conversion algorithms for a Curry-style logical framework that includes surjective pairing. Like later work by Abel and Scherer [2012], 1087 Abel and Coquand [2007] avoid the confluence problem that comes with surjective pairing by 1088 using the extensionality of the PER to model the η -laws from definitional equality. Instead, the PER 1089 model directly justifies a specific reduce-and-compare algorithm that compares the η -equivalence 1090 of β -normal forms. Similar to our development, the correctness of Coquand's algorithm is proven 1091 syntactically by showing its completeness with respect to the reduce-and-compare algorithm. The 1092 lack of a confluence result means that other variants of reduce-and-compare algorithm that may 1093 interleave β and η -reduction are not justified. Furthermore, to prove the completeness of Coquand's 1094 algorithm, Abel and Coquand start by proving it complete with respect to β -normal forms and then 1095 lift the result to include β -normalizing terms. In our development, the confluence result allows us 1096 to directly prove completeness with respect to β -normalizing terms. Finally, our proof is simpler as 1097 it only involves a logical predicate. Because convertibility in our system is modeled as untyped 1098 $\beta\eta$ -reduction, our proof is compatible with untyped definitional equality and we can add subtyping 1099 to our system without any no changes to the logical predicate. 1100

Lennon-Bertrand [2025] proves in Rocq the positive soundness, negative soundness, and termination of algorithmic conversion for a Martin-Löf type theory with one universe. The negative and positive soundness results ensure that the algorithm, upon termination, returns true precisely when its inputs are definitionally equal. This relaxation of the termination condition allows Lennon-Bertrand [2025] to formulate correctness properties for type systems that are not normalizing.

Similar to our work, Lennon-Bertrand [2025] identifies several injectivity and normalization properties that are needed to carry out a syntactic correctness proof. However, the injectivity properties required by Lennon-Bertrand [2025] differ from ours in that they are about the typed definitional equality, whereas the injectivity properties we prove to justify our algorithm are about joinability, an untyped relation.

Proving typed injectivity and normalization properties has the benefit of being compatible with 1111 singleton η -laws. However, the typed properties not only requires a more complicated logical 1112 relation, but can also cause problems when trying to separate the semantic and syntactic proofs. 1113 Lennon-Bertrand [2025] prove the *deep normalization* property, which is needed for termination, 1114 using the logical relation by Adjedj et al. [2024]. The same fundamental theorem used to show deep 1115 normalization, however, already implies the completeness of algorithmic conversion! Thus, when 1116 total correctness is concerned, the proof development by Lennon-Bertrand [2025], composed with 1117 the logical relation by Adjedj et al. [2024], does not completely decouple the logical relation from 1118 the syntactic proof about the termination of the algorithm. As part of future work, Lennon-Bertrand 1119 [2025] suggests alternative ways of showing neutral injectivity or termination that do not require 1120 the detour to the completeness result of algorithmic conversion, such as the domain theoretic 1121 method by Coquand and Huber [2019]. 1122

1124 6.2 Mechanized Metatheory for Dependent Types

1123

All of our results have been mechanized using the Rocq proof assistant [Team 2024], leading to a preference for syntactic reasoning, which is convenient and well supported by this tool. We leverage the rich ecosystem of Rocq to aid our proof development; we use CoqHammer [Czajka
and Kaliszyk 2018] to automate our proof scripts, and Autosubst [Dapprich and Dudenhefner 2021]
to generate and reason about the de Bruijn representation of our syntax. However, we are not the

first to use a proof assistant to reason about dependent type theories and we draw on prior efforts.
 Barras [1996] uses the Rocq (Coq) proof assistant to mechanize the strong normalization the

¹¹³³ Calculus of Constructions (CoC) and extract a decidable type checker. CoC, as an instance of Baren-¹¹³⁴ dregt's Pure Type Systems (PTS) [Barendregt 1991], uses untyped β -equivalence as its definitional ¹¹³⁵ equality. The conversion algorithm, which compares β -normal forms of its input, is thus directly ¹¹³⁶ justified by the confluence of β -reduction.

Liu et al. [2025] mechanize in Rocq the decidability of conversion for DCOI, a dependently typed language that supports proof-irrelevance. The definitional equality of DCOI replaces α -equivalence with a notion of syntactic indistinguishability, which ignores components of the term that cannot be distinguished by a specified observer level. The conversion algorithm checks the syntactic indistinguishability of the β -normal forms of its inputs. They prove the correctness of the algorithm syntactically by composing normalization, β -confluence, and simulation properties which relate β -reduction and syntactic indistinguishability.

Our proof shares a similar architecture to that found in Barras [1996] and Liu et al. [2025]. Both of these works use confluence-based approach to show the decidability of algorithmic conversion. This paper shows that the same approach works for systems with η -laws for functions and pairs.

Other projects that mechanize results about dependent-type theory use methods that are less related to this work.

Abel et al. [2017] use Agda to show the decidability of type conversion for a dependent type 1149 theory with one fixed universe. The algorithm used to decide type conversion is type directed and is 1150 similar to the one found in Harper and Pfenning [2005]. This work, like Harper and Pfenning [2005] 1151 and Abel and Scherer [2012], uses a Kripke-style logical relation to show the total correctness of 1152 the algorithm. To avoid having to define two separate logical relations, one for soundness and 1153 one for completeness, Abel et al. parameterize their logical relation by a relation called generic 1154 equality. A generic equality consists of an equality between types, an equality between terms, and 1155 another equality specifically for neutral terms. They then prove that the fundamental theorem 1156 holds for generic equalities that satisfy an interface of required properties. Thus, the soundness 1157 and completeness proofs now require proving that different equalities satisfy the interface, and the 1158 fundamental theorem only needs to be proven once. Adjedj et al. [2024] adopt the same technique 1159 of a parameterized logical relation in Rocq to not only show the decidability of type conversion, 1160 but also show the decidability of type checking through a bidirectional type checker. 1161

Wieczorek and Biernacki [2018] mechanize the correctness of a normalization by evaluation 1162 (NbE) algorithm for a dependent type theory with one universe in Rocq, following the pen and 1163 paper proof by Abel [2013]. Similarly, Hu et al. [2023] mechanize NbE for a modal dependent type 1164 theory with a cumulative universe in Agda. Building on the proof development by Hu et al. [2023], 1165 Jang et al. [2025] mechanize the correctness of an NbE algorithm for a more expressive dependent 1166 type theory that also supports subtyping, similar to $\lambda^{\Pi,\Sigma,U_i,\mathbb{N}}$. Jang et al. [2025], like Adjedj et al. 1167 [2024], also include a decidable type checker on top of the decidable conversion algorithm. The 1168 subtyping relation in Jang et al. [2025] is covariant on functions, whereas our subtyping relation is 1169 contravariant on functions. 1170

Altenkirch and Kaposi [2016] mechanize the correctness of NbE for a dependent type theory with a single universe in Agda. Instead of working on explicit syntax, they employ an algebraic representation of the type system as categories with families [Hofmann and Hofmann 1997], encoded as a quotient inductive type [Altenkirch et al. 2018], where terms, which are always well-typed, are quotiented by definitional equality. Unlike the proofs we have discussed so far

(including our own), Altenkirch and Kaposi [2016] use a *proof-relevant* logical predicate, where the evidence that a term is reducible is not an irrelevant proof object, but a computationally relevant structure. The proof-relevance of the logical predicate sidesteps the need for an extensional PER model to justify the η -laws [Coquand 2019]. In our work, we show that a *proof-irrelevant* logical predicate can also help avoid an extensional model by leveraging the confluence of $\beta\eta$ -reduction. Barras [2012] axiomatizes set theory in Rocq, and then models Calculus of Inductive of Constructions (CIC) using the formalized set theory. The axiomatization of set theory requires adding

additional axioms to Rocq, but also allows Barras to model the universe hierarchy with an impredicative Prop sort, which cannot be inductively defined.
Sozeau et al. [2019] prove the correctness of a type checker for a subset of Rocq that contains only

Sozeau et al. [2019] prove the correctness of a type checker for a subset of Rocq that contains only β -rules. Due to Gödel's incompleteness theorem, Sozeau et al. [2019] formulate their correctness result syntactically by assuming strong normalization.

1189 Kravchuk-Kirilyuk et al. [2020] extend the core language for Dependent Haskell [Weirich et al. 1190 2017] with function η -laws in Rocq. The decidability of conversion for System DC, an annotated 1191 version of the core calculus, has a direct syntactic proof since the type checker only needs to validate 1192 traces of reduction sequences that are part of the term syntax. The confluence of $\beta\eta$ -equivalence is 1193 used to justify the consistency its extensionally flavored equational theory.

1195 7 Conclusion and Future Work

In this work, we show the decidability of type conversion for an expressive dependent type 1196 theory with function and pair η -laws. To show the confluence of untyped $\beta\eta$ -reduction, we use 1197 the inductively defined strongly normalizing terms of Van Raamsdonk and Severi [1995]. From 1198 1199 confluence and the normalization of well-typed terms, we derive a syntactic proof of the total correctness of Coquand's algorithm (extended with subtyping) and of a reduce-and-compare 1200 algorithm. Our object language $(\lambda^{\Pi,\Sigma,U_i,\mathbb{N}})$ is expressive and demonstrates that the features of modern 1201 proof assistants, such as subtyping, large eliminations, and inductive datatypes, are compatible 1202 with this approach. We have mechanized all of our proofs using the Rocq theorem prover. 1203

Our proof method is novel as it uses a syntactic confluence proof to directly model typed convertibility as untyped $\beta\eta$ -reduction. Compared to existing proof methods, our confluence-based approach is compatible with untyped convertibility and modularly combines a minimal semantic proof of normalization with a syntactic proof of decidability. The modularity gives us the option to carry out the bulk of the decidability proof in a weak metatheory and the ability to explore different algorithms with minimal changes to the whole development.

In future work, we hope to leverage the modularity of our proof development to prove the correctness of other type conversion algorithms through syntactic means. For example, we could replace the iterated weak-head β -reduction in Coquand's algorithm with an efficient untyped NbE algorithm, which can be proven correct syntactically through confluence [Grégoire and Leroy 2002]. We also would like to apply our techniques to systems, such as ICC [Barras and Bernardo 2008; Miquel 2001] and DCOI [Liu et al. 2024b], that require untyped equality judgments and so have previously not included η -laws for dependent pairs.

References

1219Andreas Abel. 2013. Normalization by evaluation: Dependent types and impredicativity. Ph.D. Dissertation. Ludwig-1220Maximilians-Universität München.

Andreas Abel and Thierry Coquand. 2007. Untyped algorithmic equality for Martin-Löf's logical framework with surjective pairs. *Fundamenta Informaticae* 77, 4 (2007), 345–395.
 Line Abel American Construction of the Const

Andreas Abel, Thierry Coquand, and Peter Dybjer. 2008. Verifying a Semantic βη-Conversion Test for Martin-Löf Type
 Theory. In *Mathematics of Program Construction (Lecture Notes in Computer Science)*, Philippe Audebaud and Christine
 Paulin-Mohring (Eds.). Springer, Berlin, Heidelberg, 29–56. doi:10.1007/978-3-540-70594-9_4

1225

1217

1218

- Andreas Abel, Joakim Öhman, and Andrea Vezzosi. 2017. Decidability of Conversion for Type Theory in Type Theory. Proc.
 ACM Program. Lang. 2, POPL, Article 23 (Dec. 2017), 29 pages. doi:10.1145/3158111
- Andreas Abel and Gabriel Scherer. 2012. On Irrelevance and Algorithmic Equality in Predicative Type Theory. Logical Methods in Computer Science 8, 1 (2012), 1–36. doi:10.2168/lmcs-8(1:29)2012
- Robin Adams. 2006. Pure type systems with judgemental equality. Journal of Functional Programming 16, 2 (2006), 219–246.
- Arthur Adjedj, Meven Lennon-Bertrand, Kenji Maillard, Pierre-Marie Pédrot, and Loïc Pujet. 2024. Martin-Löf à la Coq. In
 Proceedings of the 13th ACM SIGPLAN International Conference on Certified Programs and Proofs (London, UK) (*CPP 2024*).
 Association for Computing Machinery, New York, NY, USA, 230–245. doi:10.1145/3636501.3636951
- Thorsten Altenkirch, Paolo Capriotti, Gabe Dijkstra, Nicolai Kraus, and Fredrik Nordvall Forsberg. 2018. Quotient inductiveinductive types. In *International Conference on Foundations of Software Science and Computation Structures*. Springer International Publishing Cham, 293–310.
- Thorsten Altenkirch and Ambrus Kaposi. 2016. Normalisation by Evaluation for Dependent Types. (2016), 16 pages.
 doi:10.4230/LIPICS.FSCD.2016.6 Artwork Size: 16 pages Medium: application/pdf Publisher: Schloss Dagstuhl Leibniz-Zentrum fuer Informatik GmbH, Wadern/Saarbruecken, Germany.
- Henk Barendregt. 1991. Introduction to generalized type systems. *Journal of Functional Programming* 1, 2 (1991), 462–490.
 doi:10.1017/s0956796800020025
- Henk P. Barendregt. 1993. Lambda Calculi with Types. Oxford University Press, Inc., USA, 117–309.
- ¹²⁴⁰ Bruno Barras. 1996. *Coq en coq.* Ph. D. Dissertation. INRIA.
- Bruno Barras. 2012. Semantical investigations in intuitionistic set theory and type theories with inductive families. Thèse
 l'habilitation à diriger des recherches, Université Paris 7 Denis Diderot.
- Bruno Barras and Bruno Bernardo. 2008. The Implicit Calculus of Constructions as a Programming Language with Dependent Types. In *Foundations of Software Science and Computational Structures*, Roberto Amadio (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 365–379. doi:10.1007/978-3-540-78499-9_26
- Ulrich Berger, Matthias Eberl, and Helmut Schwichtenberg. 1998. Normalization by evaluation. Prospects for Hardware
 Foundations: ESPRIT Working Group 8533 NADA—New Hardware Design Methods Survey Chapters (1998), 117–137.
- 1247Ana Bove and Venanzio Capretta. 2005. Modelling general recursion in type theory. Mathematical Structures in Computer1248Science 15, 4 (2005), 671–708. doi:10.1017/S0960129505004822
- Pritam Choudhury, Harley Eades III, and Stephanie Weirich. 2022. A Dependent Dependency Calculus. In *Programming Languages and Systems, ESOP 2022 (Lecture Notes in Computer Science, Vol. 13240)*, Ilya Sergey (Ed.). Springer International Publishing, Cham, 403–430. doi:10.1007/978-3-030-99336-8_15 Artifact available.
- 1251 Thierry Coquand. 1991. An algorithm for testing conversion in type theory. Logical frameworks 1 (1991), 255–279.
- 1252Thierry Coquand. 2019. Canonicity and normalization for dependent type theory. Theoretical Computer Science 777 (July12532019), 184–191. doi:10.1016/j.tcs.2019.01.015
- Thierry Coquand and Simon Huber. 2019. An Adequacy Theorem for Dependent Type Theory. *Theory of Computing Systems* 63, 4 (May 2019), 647–665. doi:10.1007/s00224-018-9879-9
- Lukasz Czajka and Cezary Kaliszyk. 2018. Hammer for Coq: Automation for dependent type theory. *Journal of automated reasoning* 61 (2018), 423–453.
- Adrian Dapprich and Andrej Dudenhefner. 2021. Generating Infrastructural Code for Terms with Binders using MetaCoq
 and OCaml. Bachelor thesis, Saarland University (2021).
- Herman Geuvers. 1994. A short and flexible proof of strong normalization for the calculus of constructions. In *International Workshop on Types for Proofs and Programs*. Springer, Berlin, Heidelberg, 14–38. doi:10.1007/3-540-60579-7_2
 Ian Horman Geuvers. 1903. Logics and two systems. [SI: cal.]
 - Jan Herman Geuvers. 1993. Logics and type systems. [Sl: sn].
- 1261 Jean-Yves Girard, Paul Taylor, and Yves Lafont. 1989. Proofs and types. Vol. 7. Cambridge University Press, Cambridge.
- Healfdene Goguen. 2005. Justifying Algorithms for βη-Conversion. In Foundations of Software Science and Computational Structures, David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Dough Tygar, Moshe Y.
- Vardi, Gerhard Weikum, and Vladimiro Sassone (Eds.). Vol. 3441. Springer Berlin Heidelberg, Berlin, Heidelberg, 410–424.
 doi:10.1007/978-3-540-31982-5_26 Series Title: Lecture Notes in Computer Science.
- Benjamin Grégoire and Xavier Leroy. 2002. A compiled implementation of strong reduction. In *Proceedings of the seventh* ACM SIGPLAN international conference on Functional programming. 235–246.
- Robert Harper and Frank Pfenning. 2005. On equivalence and canonical forms in the LF type theory. ACM Transactions on Computational Logic (TOCL) 6, 1 (2005), 61–101.
- Martin Hofmann and Martin Hofmann. 1997. Syntax and semantics of dependent types. *Extensional Constructs in Intensional Type Theory* (1997), 13–54.
- Jason Z. S. Hu, Junyoung Jang, and Brigitte Pientka. 2023. Normalization by evaluation for modal dependent type theory.
 Journal of Functional Programming 33 (2023), e7. doi:10.1017/S0956796823000060
- 1273

Junyoung Jang, Jason ZS Hu, Antoine Gaulin, and Brigitte Pientka. 2025. McTT: Building A Correct-By-Construction Proof
 Checker For Martin-Löf Type Theory. (2025).

27

- Felix Joachimski and Ralph Matthes. 2003. Short proofs of normalization for the simply- typed λ-calculus, permutative conversions and Gödel's T. Archive for Mathematical Logic 42, 1 (Jan. 2003), 59–87. doi:10.1007/s00153-002-0156-9
- Jan Willem Klop. 1980. Combinatory Reduction Systems. Mathematisch centrum, Amsterdam.
- J. W. Klop and R. C. de Vrijer. 1989. Unique normal forms for lambda calculus with surjective pairing. *Information and Computation* 80, 2 (Feb. 1989), 97–113. doi:10.1016/0890-5401(89)90014-X
- Anastasiya Kravchuk-Kirilyuk, Antoine Voizard, and Stephanie Weirich. 2020. Eta-Equivalence in Core Dependent Haskell.
 Leibniz international proceedings in informatics 175 (2020).
- 1283 Meven Lennon-Bertrand. 2022. À bas l' η —Coq's troublesome η -conversion. The first Workshop on the Implementation of Type Systems (WITS). https://www. meven. ac/documents/22-WITS-abstract. pdf.
- Meven Lennon-Bertrand. 2025. What Does It Take to Certify a Conversion Checker?. In 10th International Conference on Formal Structures for Computation and Deduction (FSCD 2025) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 337), Maribel Fernández (Ed.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 27:1–27:23. doi:10.4230/LIPIcs.FSCD.2025.27
- 1288 Yiyun Liu, Jonathan Chan, Jessica Shi, and Stephanie Weirich. 2024b. Internalizing Indistinguishability with Dependent Types. *Proc. ACM Program. Lang.* 8, POPL, Article 44 (Jan. 2024), 28 pages. doi:10.1145/3632886
- Yiyun Liu, Jonathan Chan, and Stephanie Weirich. 2024a. Functional Pearl: Short and Mechanized Logical Relation for
 Dependent Type Theories. (2024).
- 1291Yiyun Liu, Jonathan Chan, and Stephanie Weirich. 2025. Consistency of a Dependent Calculus of Indistinguishability. Proc.1292ACM Program. Lang. 9, POPL, Article 7 (Jan. 2025), 27 pages. doi:10.1145/3704843
- Zhaohui Luo. 1990. An extended calculus of constructions. Ph. D. Dissertation. University of Edinburgh.
- Alexandre Miquel. 2001. The Implicit Calculus of Constructions Extending Pure Type Systems with an Intersection Type
 Binder and Subtyping. In *Typed Lambda Calculi and Applications*, Samson Abramsky (Ed.). Springer Berlin Heidelberg,
 Berlin, Heidelberg, 344–359. doi:10.1007/3-540-45413-6_27
- 1296Alexandre Miquel and Benjamin Werner. 2003. The Not So Simple Proof-Irrelevant Model of CC. In Types for Proofs and1297Programs, Herman Geuvers and Freek Wiedijk (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 240–258.
- Vincent Siles and Hugo Herbelin. 2012. Pure type system conversion is always typable. *Journal of Functional Programming* 22, 2 (2012), 153–180. doi:10.1017/S0956796812000044
- Matthieu Sozeau, Simon Boulier, Yannick Forster, Nicolas Tabareau, and Théo Winterhalter. 2019. Coq Coq correct!
 verification of type checking and erasure for Coq, in Coq. *Proc. ACM Program. Lang.* 4, POPL, Article 8 (Dec. 2019),
 28 pages. doi:10.1145/3371076
- Jonathan Sterling and Carlo Angiuli. 2021. Normalization for cubical type theory. In 2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS). IEEE, 1–15.
- Masako Takahashi. 1995. Parallel Reductions in λ -Calculus. Information and Computation 118, 1 (1995), 120–127. doi:10. 1006/inco.1995.1057
- 1305 The Coq Development Team. 2024. The Coq Proof Assistant. doi:10.5281/zenodo.14542673
- 1306Amin Timany and Matthieu Sozeau. 2017. Consistency of the Predicative Calculus of Cumulative Inductive Constructions1307(pCuIC). Research Report RR-9105. KU Leuven, Belgium ; Inria Paris. 30 pages. https://inria.hal.science/hal-01615123
- Femke Van Raamsdonk and PG Severi. 1995. On normalisation. (1995).
- 1000Stephanie Weirich, Antoine Voizard, Pedro Henrique Avezedo de Amorim, and Richard A. Eisenberg. 2017. A Specification1309for Dependent Types in Haskell. Proc. ACM Program. Lang. 1, ICFP, Article 31 (Aug. 2017), 29 pages. doi:10.1145/3110275
- Paweł Wieczorek and Dariusz Biernacki. 2018. A Coq Formalization of Normalization by Evaluation for Martin-Löf Type
 Theory. In *Proceedings of the 7th ACM SIGPLAN International Conference on Certified Programs and Proofs* (Los Angeles, CA, USA) (*CPP 2018*). Association for Computing Machinery, New York, NY, USA, 266–279. doi:10.1145/3167091
- 1313
- 1314
- 1315 1316