A Reasonably Exceptional Type Theory

PIERRE-MARIE PÉDROT, Inria, France

NICOLAS TABAREAU, Inria, France

HANS JACOB FEHRMANN, University of Chile, Chile

ÉRIC TANTER, University of Chile, Chile and Inria Paris, France

Traditional approaches to compensate for the lack of exceptions in type theories for proof assistants have severe drawbacks from both a programming and a reasoning perspective. Pédrot and Tabareau recently extended the Calculus of Inductive Constructions (CIC) with exceptions. The new exceptional type theory is interpreted by a translation into CIC, covering full dependent elimination, decidable type-checking and canonicity. However, the exceptional theory is inconsistent as a logical system. To recover consistency, Pédrot and Tabareau propose an additional translation that uses parametricity to enforce that all exceptions are caught locally. While this enforcement brings logical expressivity gains over CIC, it completely prevents reasoning about exceptional programs such as partial functions. This work addresses the dilemma between exceptions and consistency in a more flexible manner, with the *Reasonably Exceptional Type Theory* (RETT). RETT is structured in three layers: (a) the *exceptional* layer, in which all terms can raise exceptions; (b) the *mediation* layer, in which exceptional terms must be provably parametric; (c) the *pure* layer, in which terms are non-exceptional, but can refer to exceptional terms. We present the general theory of RETT, where each layer is realized by a predicative hierarchy of universes, and develop an instance of RETT in Coq: the impure layer corresponds to the predicative universe hierarchy, the pure layer is realized by the impredicative universe of propositions, and the mediation layer is reified via a parametricity type class. RETT is the first full dependent type theory to support consistent reasoning about exceptional terms, and the CoqRETT plugin readily brings this ability to Coq programmers.

CCS Concepts: • **Theory of computation** → **Type theory**; **Functional constructs**; *Program verification*.

Additional Key Words and Phrases: type theory, effects, exceptions, parametricity

ACM Reference Format:

Pierre-Marie Pédrot, Nicolas Tabareau, Hans Jacob Fehrmann, and Éric Tanter. 2019. A Reasonably Exceptional Type Theory. *Proc. ACM Program. Lang.* 3, ICFP, Article 108 (August 2019), 29 pages. <https://doi.org/10.1145/3341712>

1 FAILURE IN TYPE THEORY

The absolute purity of type theories like the Calculus of Constructions [Coquand and Huet 1988] is both a blessing and a curse. A blessing because purity implies consistency of the internal logic, thereby validating their use as foundations of proof assistants such as Coq [The Coq Development Team 2019] and Agda [Norell 2009], within which one can express and prove interesting

*This work is partially funded by CONICYT FONDECYT Regular Project 1190058, by CONICYT REDES Project 170067, by the European Research Council under ERC Starting Grant CoqHoTT (637339) and ERC Starting Grant SECOMP (715753) and by Inria Équipe Associée GECO.

Authors' addresses: Pierre-Marie Pédrot, Inria, Gallinette Project-Team, Nantes, France; Nicolas Tabareau, Inria, Gallinette Project-Team, Nantes, France; Hans Jacob Fehrmann, University of Chile, Computer Science Department (DCC), Santiago, Chile; Éric Tanter, University of Chile, Computer Science Department (DCC), Santiago, Chile, Inria Paris, France.

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

© 2019 Copyright held by the owner/author(s).

2475-1421/2019/8-ART108

<https://doi.org/10.1145/3341712>

mathematical results, including about programs and programming languages. A curse because the lack of a basic effect like failure makes the use of these theories in practical scenarios, in particular in their dual use as functional programming languages, cumbersome at best.

The Failure Problem. As a matter of fact, many common situations would benefit from a convenient way to deal with partiality or failure, like exceptions in mainstream programming languages. We will call this the *failure problem*. Traditional solutions to the failure problem are monadic programming, default values, and axioms; each of which has severe drawbacks as discussed next.

Monadic Programming. The standard approach to the failure problem in functional programming is to use the option (or exception) monad, in which values are tagged explicitly to indicate whether they denote a success or a failure. For instance, the head function on lists can be given type $\Pi A : \square. \text{list } A \rightarrow \text{option } A$, where $\square$ is the universe (a.k.a. kind) of types. This approach is notoriously contagious, widely imposing a monadic style of programming. More problematic, while it can be used without too much pain in a non-dependently-typed setting like e.g. in Haskell, the monadic approach does not scale well to dependent types.

For instance, if a function f returns an optional value, then even a simple dependent property like $\forall x. x > 0 \rightarrow f\ x > 0$ needs to be stated with type-level pattern matching, such as: $\forall x. x > 0 \rightarrow \text{match } f\ x \text{ with } \text{Some } y \rightarrow y > 0 \mid \text{None} \rightarrow ???$. In addition to quickly becoming untractable, this technique is also not systematic: what type should one put in place of $???$ in the failure branch? For positive occurrences, a top type seems useful, but dually, for negative occurrences, a bottom type should be considered to rule out exceptional cases.

Default Values. Due to these issues, many libraries tend to favor the use of default values over the exception monad. Default values preserve simple signatures, but for polymorphic functions, coming up with a default value is tricky, requiring either the use of type class resolution to automatically infer default values for certain (inhabited) types—an approach used for instance in *hs-to-coq* [Breitner et al. 2018]—or changing the signature of functions to require as first argument a default value to return in case of failure—an approach favored in the Mathematical Components library [Mahboubi and Tassi 2008], where e.g. $\text{head} : \Pi A : \square. A \rightarrow \text{list } A \rightarrow A$.

A major drawback of using default values is the potential confusion with normal values: if head returns the default value, is it because the list was empty, or because the default value was actually the head of the list? Consequently, reasoning about such artificially total functions is compromised. Likewise, how to state that “the tail of a non-empty list does not fall through the problematic branch”? The property

$$\text{length } l > 0 \rightarrow \text{tail } d\ l \neq d$$

is not true in general: it does not hold if d is the empty list and l has just one element.

Axiomatic Approach. To avoid imposing a monadic style while avoiding confusion of values, another approach is to use axioms to denote exceptions, as explored by Tanter and Tabareau [2015] in their cast framework for subset types. The problem of axioms is that they have no computational content, therefore raising an exception materializes as a stuck term; it is impossible to catch such axiomatic exceptions and to reason about potential failing terms. For instance, if tail uses an axiom error when applied to an empty list, the property

$$\text{length } l > 0 \rightarrow \text{tail } l \neq \text{error}$$

is not provable, because one is not allowed discriminate the axiom from pure terms.

Exceptional Type Theory. Recently, Pédro and Tabareau [2018] developed an extension of CIC with exceptions, interpreted by a translation into CIC, and implemented in Coq as a plugin. The

Exceptional Type Theory (ETT) includes a function $\text{fail} : \Pi A : \square. A$ that throws an exception at any type (for readability, we omit the type argument in the remainder of this section). This function enjoys the computational behavior that one would expect, namely that the exception escapes from contexts that evaluate it.

This solves the failure problem in a straightforward way. ETT makes it possible to define *head* and *tail* functions that raise exceptions when applied to the empty list, without polluting their type signatures, and to prove that

$$\vdash_{\text{ETT}} \text{length } l > 0 \rightarrow \text{tail } l \neq \text{fail}.$$

The good news is that ETT is computationally relevant, that is, programs reduce to normal forms and the equational theory is not degenerate. As such, it can be used as a dependently-typed programming language with exceptions.

The flip side is that ETT is inconsistent as a logic. Just as in any programming language featuring exceptions, it is indeed possible to inhabit any type, and thus any property, by raising an exception. In particular, ETT *also* allows one to prove the paradoxical fact that

$$\vdash_{\text{ETT}} \text{length } l > 0 \rightarrow \text{tail } l = \text{fail}.$$

Peeking at the proofs of those two properties reveals that they are not anywhere near being equally valid, though. The first one is correct in the sense that it is only using the exception-free fragment of ETT, while the second one is a blatant lie, as executing it would lead to an immediate dynamic failure.

By defining a subset of *valid* proofs, it is possible to recover logical consistency. [Pédrot and Tabareau \[2018\]](#) give a second interpretation of CIC as a restriction of ETT, using parametricity [\[Bernardy and Lasson 2011\]](#) to force all exceptions to be locally handled. This approach is useful to extend the logical expressivity of CIC with a kind of backtracking-based reasoning. Unfortunately, the restriction is too strong and is not applicable to the programming setting considered here. One cannot define exception-raising functions like *head* or *tail* anymore, because by construction they do not satisfy the validity criterion.

Consistent Reasoning about Exceptional Programs. The contribution of this paper is to present a new type theory, dubbed the Reasonably Exceptional Type Theory (RETT), which supports consistent reasoning about exceptional programs. The core feature of RETT that makes this possible is a universe-based separation between consistent proofs and effectful programs. This split is embodied by the existence of parallel hierarchies of safe vs. unsafe types that are allowed to interact in a principled way.

We implement CoQRETT, a fragment of RETT in CoQ as a plugin. Seizing their similarity of purpose, CoQRETT piggybacks on CoQ's Prop-Type classification to separate consistent logical reasoning (in Prop) from effectful programming (in Type). We convey a foretaste of CoQRETT in the paragraphs below.

In CoQRETT, we can define *head* and *tail* as functions that raise exceptions when applied to empty lists, as those terms live in the Type hierarchy. We can then prove as expected

$$\text{length } l > 0 \rightarrow \text{tail } l \neq \text{fail}.$$

Contrarily to ETT, in CoQRETT the paradoxical proposition $\text{length } l > 0 \rightarrow \text{tail } l = \text{fail}$ does not hold, because equality lives in Prop, forbidding inconsistent reasoning. Similarly and somehow counter-intuitively, in CoQRETT we cannot prove that

$$\Pi n : \mathbb{N}. n \geq 0.$$

assuming $\geq$ is defined in the usual way. The reason is that exceptions also inhabit the Type-dwelling type of natural numbers $\mathbb{N}$, while $\geq$ only mentions pure integers. In order to be able to reason

about terms that are actually pure, we need to introduce a parametricity predicate `param`, realized using a Coq type class. This way, we can prove

$$\prod n : \mathbb{N}. \text{param } n \rightarrow n \geq 0.$$

This selective and explicit approach to parametricity is the key to allow both consistent reasoning and exceptional terms to coexist.

Pédrot and Tabareau [2018] observe that exceptions in type theory are naturally call-by-name exceptions. This means for instance that there are multiple levels at which “being a list of natural numbers” can interact with failure: the whole list can be an exception, the spine of the list can be a proper structure but it can have exceptions as elements (i.e. fake natural numbers), or the list can be a pure, deeply parametric list.

To illustrate, consider a property of an exception-raising head function, namely that it does not fail when applied to non-empty lists

$$\text{length } l > 0 \rightarrow \text{head } l \neq \text{fail}.$$

Stated in this way, this property is in fact false. Is it because l itself could be an exception? No. In fact, we can prove that $\text{length } l > 0 \rightarrow l \neq \text{fail}$, because if $l = \text{fail}$, then $\text{length } l$ is convertible to `fail`; but the proposition $\text{fail} > 0$ cannot be proven in `Prop`, which is consistent. In other words, $\text{length } l > 0 \rightarrow \text{param } l$.

But even though l is a “real” list, its elements might not be. In particular, the first element might be an exception, thereby negating the property above. In order to properly state the property, we therefore need a deep version of the parametricity predicate. We can then prove that

$$\text{param}_{\text{deep}} l \rightarrow \text{length } l > 0 \rightarrow \text{head } l \neq \text{fail}.$$

In brief, CoqRETT allows programmers to use exceptions in their programs, and to consistently reason about them at the required level of granularity, accounting for potential failures extrinsically and when needed. We come back to these examples in Section 5.3, presenting in detail their statements and proofs in CoqRETT.

A Reasonably Exceptional Type Theory. CoqRETT represents the practical contribution of this work. However, as the illustration above reveals, CoqRETT relies in an essential way on the parametricity predicate `param`, which is realized through a type class. This is problematic from a foundational point of view, because type classes and ad hoc polymorphism cannot be directly accounted for in a type-theoretic setting.

To put consistent reasoning about exceptional terms on a solid type theoretic footing, we propose the Reasonably Exceptional Type Theory (RETT). RETT features three separate universe hierarchies, which can be thought of as adjacent layers: one exceptional, one pure, and in between a mediation layer in which parametricity is realized (Section 3). Modalities, defined as functions, coordinate the interplay between these layers. We give a syntactic model of RETT by translation into CIC, interpreting each layer and modality in a specific way (Section 4). This translation allows us to prove the metatheoretical properties of RETT. CoqRETT is then formally justified by considering a fragment of RETT that is implementable in Coq (Section 5) without having to modify its kernel. Coq being restricted to two universe hierarchies, we map the exceptional layer to the predicative hierarchy (Type), and the pure layer to the impredicative universe (Prop); type classes are then an implementation technique to reify the parametricity predicate from the (missing) mediation layer.

The implementation of the CoqRETT plugin and the examples discussed here are provided in supplementary material; they have been tested in Coq 8.8.

2 BACKGROUND: EXCEPTIONAL TRANSLATION AND PARAMETRICITY

We first provide a quick introduction to the key technical ideas of the Exceptional Type Theory (ETT) of [Pédrot and Tabareau \[2018\]](#), on which our technical development is based. We recall both interpretations of ETT: the standard exceptional translation, which yields a logically inconsistent theory; and the parametric exceptional translation, which recovers consistency through parametricity, at the expense of expressiveness.

2.1 Exceptional Translation

As mentioned in the introduction, ETT is an extension of CIC with exceptions. ETT includes an exception type $E : \square$ and a function $\text{raise} : \Pi A : \square. E \rightarrow A$ to raise exceptions at any type A . ETT is justified by a syntactic translation into CIC, denoted $[M]$ for any ETT term M , which is a simplification of the weaning translation of [Pédrot and Tabareau \[2017\]](#). Intuitively, a type A in ETT is interpreted as a pair of a type A in CIC together with a default function $A_\emptyset : \mathbb{E} \rightarrow A$ specifying how to interpret failure on this type. Here, $\mathbb{E}$ is the CIC representation type of the source exception type E .

Because the universe of types is itself a type, one needs to define a representation for types that can raise exceptions. This can be done with the following inductive type:

$$\begin{aligned} \text{Ind type}_i : \square_{i+1} := \\ & | \text{TypeVal}_i : \Pi A : \square_i. (\mathbb{E} \rightarrow A) \rightarrow \text{type}_i \\ & | \text{TypeErr}_i : \mathbb{E} \rightarrow \text{type}_i \end{aligned}$$

The constructor TypeVal_i constructs a type_i from a type and a default function on this type. The constructor TypeErr_i represents the default function at the level of type_i . The exceptional translation uses a term $\text{El}_i : \text{type}_i \rightarrow \square_i$ to recover the underlying type from an inhabitant of type_i , and a term $\text{Err}_i : \Pi A : \text{type}_i. \mathbb{E} \rightarrow \text{El}_i A$ to lift the default function to this underlying type.

The translation of an ETT universe is therefore a value of the above inductive type:

$$[\square_i] := \text{TypeVal}_{i+1} \text{ type}_i \text{ TypeErr}_i$$

The ETT exception type E is mapped to $\mathbb{E}$ together with the identity as default function:

$$[E] := \text{TypeVal } \mathbb{E} (\lambda e : \mathbb{E}. e)$$

and the function raise raises the provided exception at any type as:

$$[\text{raise}] := \lambda(A : \text{type})(e : \mathbb{E}). \text{Err } A e$$

For inductive types, the translation freely adds an additional constructor, similarly to TypeErr for the universe. For instance, the translation of the inductive type of booleans $\mathbb{B}$ is

$$[\mathbb{B}] := \text{TypeVal } \mathbb{B}^\bullet \mathbb{B}_\emptyset$$

where $\mathbb{B}^\bullet$ is an inductive type with three constructors: $\text{true}^\bullet$, $\text{false}^\bullet$ and an extra default constructor $\mathbb{B}_\emptyset : \mathbb{E} \rightarrow \mathbb{B}^\bullet$.

Observe that this treatment of inductive types means that the empty type of ETT is translated to an inductive with one (default) constructor. Therefore the empty type is inhabited as soon as the target exception type $\mathbb{E}$ is. In fact, any proof of the empty type is an exception, which means that one can use exceptions to prove any result. ETT is inconsistent as a logic.

2.2 Exceptional Parametric Translation

To recover logical consistency, [Pédrot and Tabareau \[2018\]](#) give a second interpretation of ETT that uses the standard parametricity translation for type theory [[Bernardy and Lason 2011](#)] in addition to the exceptional translation.

Let us first recall the essence of the (unary) parametricity translation for type theory. While in System F, parametricity has to be stated and proved externally, the expressiveness of type theory makes it possible to internalize the parametricity argument as a translation from terms to terms. The mere fact that the translation is defined for all terms means that these terms are parametric—this property is known as the *abstraction theorem* [Reynolds 1983]. In type theory, for the universe, the parametricity translation is defined as arbitrary predicates on types, e.g. type A is translated to a predicate of type $A \rightarrow \square$. These predicates are called *parametricity predicates*, inhabited by *parametricity witnesses*. For the dependent function type $\Pi x : A. B$, the translation specifies that a *valid* input, i.e. an argument of type A together with its parametricity witness, yields a *valid* output at type B . Consequently, the translation of a lambda term is a function takes an argument and a parametricity witness, and an application is translated so as to pass the parametricity witness as extra argument.

In the context of ETT, one can simply use the parametricity translation to detect pure terms by postulating the *absence* of any parametricity witness for *raise*. In this way, given a term, if its parametricity translation is defined, then the term does not use *raise*. While this violently rules out any use of exceptions, one can manually extend the parametricity translation for a term t that uses exceptions *internally*: all that is required is to exhibit a proof that t satisfies the parametricity predicate corresponding to its type. Pédrot and Tabareau [2018] exploit this approach to show the independence of premises with a term that uses exceptions locally as a backtracking reasoning technique.

The exceptional parametric translation, denoted $[-]_e$, is the standard parametricity translation parametrized by the exceptional translation $[-]$. The exceptional parametric translation enjoys an abstraction theorem similar to standard parametricity, stated as follows:

THEOREM 2.1 (PÉDROT AND TABAREAU [2018]).

If $\vdash M : A$ and $[M]_e$ is defined, then $\vdash [M]_e : [A]_e [M]$

The condition of $[M]_e$ to be defined captures the extensibility of parametric reasoning in ETT: the translation is automatically defined for all CIC terms, and can be extended manually for terms that properly handle all their exceptions internally.

The purpose of the predicate $[A]_e$ on a type A is to forbid the use of *raise* to inhabit it. Any type A in ETT is turned into a parametricity predicate $[A]_e : \llbracket A \rrbracket \rightarrow \square$, which encodes the fact that an inhabitant of A is not allowed to generate an uncaught exception. Any occurrence of a term of the original theory used in the parametricity translation is replaced by its exceptional translation, using $[\cdot]$ or $\llbracket \cdot \rrbracket$ depending on whether it is used as a term or as a type. For instance, the translation of an application $[M N]_e$ is given by $[M]_e [N] [N]_e$ instead of just $[M]_e N [N]_e$.

The translation of the universe is given by

$$[\square]_e := \lambda A : \llbracket \square \rrbracket. \llbracket A \rrbracket \rightarrow \square_i$$

where $\llbracket A \rrbracket := E1 [A]$ is the translation of a term seen as a type (i.e., on the right-hand side of a typing judgment).

For inductive types, the default (error) constructor is always invalid, while all other constructors are valid, assuming their arguments are. For instance, the parametric translation $\mathbb{B}_e : \mathbb{B}^* \rightarrow \square$ for the inductive type $\mathbb{B}$ is an inductive type with only two constructors: $\text{true}_e : \mathbb{B}_e \text{ true}^*$ and $\text{false}_e : \mathbb{B}_e \text{ false}^*$. This means that true^* and false^* are parametric, but $\mathbb{B}_\emptyset$ is not.

As explained in the introduction, this new interpretation ensures logical consistency but rules out defining functions that let exceptions escape, let alone reasoning about such exceptional terms.

3 REASONABLY EXCEPTIONAL TYPE THEORY: DEFINITION

The Reasonably Exceptional Type Theory (RETT) supports consistent reasoning about exceptional programs by clearly separating three different universe hierarchies, and providing modalities to interoperate between them.

The (predicative) universe hierarchies of RETT are:

- the **exceptional layer**, $\Box^e$: this layer corresponds to plain ETT. It features an exception type E together with a failure function `raise`, and as such, is logically inconsistent.
- the **mediation layer**, $\Box^m$: this layer corresponds to the parametric fragment of ETT. While exceptions exist internally there, one must ensure that they are all caught before reaching toplevel. This safety discipline ensures consistency *a posteriori*.
- the **pure layer**, $\Box^p$: this layer only features the standard constructions of CIC. In particular, it does not allow raising exceptions at all.

The mediation layer supports the internalization of a parametricity predicate that classifies effectful terms that happen to be pure. All three layers are interpreted by translations to CIC, which in particular allows us to prove consistency and canonicity for the mediation and pure layers.

The interplay between these layers is threefold. First, dependent products are allowed to quantify over one layer in their domain, and another layer in their codomain (Section 3.1). Second, inductive types can be defined in all three layers, and can be eliminated into any other layer (Section 3.2). Crucially, elimination principles of inductive types depend on the source and target layers, in order to avoid compromising consistency. For instance, eliminating from $\Box^e$ to $\Box^m$ or $\Box^p$ requires explicitly accounting for potential exceptions. Third, RETT provides *modalities*—i.e. functions—connecting the mediation layer to the other two, both ways (Section 3.3). In particular, this allows RETT to feature an internal parametricity predicate that classifies effectful terms that happen to be pure (Section 3.4); this is key to allow generic reasoning about purity of exceptional terms. Finally, modalities and the parametricity predicate can be used to inject inductive types between layers (Section 3.5). This allows us to recover the standard elimination principle in the mediation layer, for impure inductive types of the exception layer, providing that the term on which we do elimination is parametric.

Why three layers? One may be surprised by the existence of three layers, rather than just two, namely one for effectful programming and one for pure reasoning. In a nutshell, this is because we cannot have a single layer that is both compatible with extensional properties (e.g. function extensionality), and suitable to define an internal parametricity predicate.

The pure layer satisfies the former but not the latter. It is merely an embedding of CIC, and thus proves the same theorems when they do not involve effectful programs. This conservativity result (Section 4.7) is very important as it allows us to backport any additional property one may want to add to CIC in the pure layer. For instance, the pure layer is compatible with functional extensionality, univalence or uniqueness of identity proofs. The disadvantage is that it does not give rise to an internal parametricity predicate, and so does not allow generic reasoning about purity of exceptional terms. This is intuitively because the pure layer is completely agnostic to exceptions. This intuition is made precise thanks to the translation presented in Section 4.5.

Dually, the mediation layer is but a logically-consistent restriction of the exceptional layer: it contains exceptions, but tamed by parametricity. It thus enables the definition of an internal parametricity predicate by using the modality from $\Box^m$ to $\Box^e$. Again, this intuition is made precise in Section 4.5. The price to pay is that the mediation layer gains impure property from the presence of internal exceptions, most notably the fact that it negates function extensionality. See Section 4.7 for more on this topic.

Syntax

$$\begin{aligned}
 s & ::= e \mid m \mid p \\
 A, B, M, N, P & ::= \Box_i^s \mid x \mid M^s N \mid \lambda x :^s A. M \mid \Pi x :^s A. B \mid \mathbf{const} \\
 \Gamma, \Delta & ::= \cdot \mid \Gamma, x :^s A
 \end{aligned}$$

Rules

$$\begin{array}{c}
 \frac{}{\vdash \cdot} \quad \frac{\Gamma \vdash A : \Box_i^s}{\vdash \Gamma, x :^s A} \quad \frac{\Gamma \vdash A : \Box_i^s}{\Gamma, x :^s A \vdash x : A} \quad \frac{\Gamma \vdash M : B \quad \Gamma \vdash A : \Box_i^s}{\Gamma, x :^s A \vdash M : B} \\
 \\
 \frac{\Gamma \vdash M : B \quad \Gamma \vdash A : \Box_i^s \quad A \equiv B}{\Gamma \vdash M : A} \quad \frac{\vdash \Gamma \quad i < j}{\Gamma \vdash \Box_i^s : \Box_j^s} \\
 \\
 \frac{\Gamma \vdash A : \Box_i^{s_1} \quad \Gamma, x :^{s_1} A \vdash B : \Box_j^{s_2}}{\Gamma \vdash \Pi x :^{s_1} A. B : \Box_{\max(i,j)}^{s_2}} \quad \frac{\Gamma, x :^{s_1} A \vdash M : B \quad \Gamma \vdash \Pi x :^{s_1} A. B : \Box_i^{s_2}}{\Gamma \vdash \lambda x :^{s_1} A. M : \Pi x :^{s_1} A. B} \\
 \\
 \frac{\Gamma \vdash M : \Pi x :^s A. B \quad \Gamma \vdash N : A}{\Gamma \vdash M^s N : B\{x := N\}}
 \end{array}$$

Constants

$$\begin{aligned}
 \mathbf{E} & : \Box_0^e \\
 \mathbf{raise} & : \Pi A : \Box_i^e. \mathbf{E} \rightarrow A
 \end{aligned}$$

Fig. 1. RETT: Syntax and typing rules

$$\begin{aligned}
 (\lambda x :^s A. M)^s N & \equiv M\{x := N\} \quad (\text{congruence rules omitted}) \\
 \mathbf{raise}^{s_2}(\Pi x :^{s_1} A. B)^e M & \equiv \lambda x :^{s_1} A. \mathbf{raise}^{s_2} B^e M
 \end{aligned}$$

Fig. 2. RETT: Conversion rules

3.1 Negative Fragment

Figure 1 presents the syntax and typing rules of RETT. Apart from the three universe hierarchies and their corresponding binder and application annotations, the syntax is standard.

The first rules are standard and apply for all hierarchies. To make layer constraints explicit, we use $\Box^s$ where s ranges over the layer identifiers e , m , and p . For instance, the universe rule specifies that each layer contains a denumerable well-founded sequence, but isolated from each other. Although the RETT syntactic model supports it, for simplicity the system featured in this paper does not feature cumulativity.

For technical reasons, we also annotate binders and applications with the layer in which the type of their argument lives, but we will often omit these annotations when they are clear from the context. Importantly, the dependent product is allowed to quantify over a type A from a universe in any layer; the layer of the dependent product type is determined by the layer of its codomain. This means that the dependent product crosscuts the three hierarchies, e.g. it is sufficient for B to live in $\Box^p$ to ensure that $\Pi x :^e A. B$ lives in $\Box^p$, even when quantifying over A in $\Box^e$.

The typing of exceptions is the same as in ETT, save for the extra precision of the universe hierarchy: the exception type $\mathbf{E}$ lives in the exceptional layer, and likewise $\mathbf{raise}$ can only raise types from $\Box^e$.

Inductive constants

$$\begin{aligned}
\mathbb{B}^e & : \square_i^e \\
\text{true}^e & : \mathbb{B}^e \\
\text{false}^e & : \mathbb{B}^e \\
\text{rec}_{\mathbb{B}^e} & : \Pi P : \mathbb{B}^e \rightarrow \square_i^e. P \text{ true}^e \rightarrow P \text{ false}^e \rightarrow \Pi b : \mathbb{B}^e. P b \\
\mathbb{B}^m & : \square_i^m \\
\text{true}^m & : \mathbb{B}^m \\
\text{false}^m & : \mathbb{B}^m \\
\text{rec}_{\mathbb{B}^m} & : \Pi P : \mathbb{B}^m \rightarrow \square_i^m. P \text{ true}^m \rightarrow P \text{ false}^m \rightarrow \Pi b : \mathbb{B}^m. P b \\
\text{list} & : \square_i^m \rightarrow \square_i^m \\
\text{nil} & : \Pi A : \square_i^m. \text{list } A \\
\text{cons} & : \Pi A : \square_i^m. A \rightarrow \text{list } A \rightarrow \text{list } A \\
\text{rec}_{\text{list}} & : \Pi(A : \square_i^m)(P : \text{list } A \rightarrow \square_i^m). \\
& \quad P(\text{nil } A) \rightarrow (\Pi(x : A)(l : \text{list } A). P l \rightarrow P(\text{cons } A x l)) \rightarrow \Pi l : \text{list } A. P l \\
\text{eq} & : \Pi A : \square_i^p. A \rightarrow A \rightarrow \square_i^p \\
\text{refl} & : \Pi(A : \square_i^p)(x : A). \text{eq } A x x \\
\text{rec}_{\text{eq}} & : \Pi(A : \square_i^p)(x : A)(P : \Pi y : A. \text{eq } A x y \rightarrow \square_i^p). P x (\text{refl } A x) \rightarrow \Pi(y : A)(e : \text{eq } A x y). P y e
\end{aligned}$$
Conversion rules

$$\begin{aligned}
\text{rec}_{\mathbb{B}^m} P P_t P_f \text{true}^m & \equiv P_t & \text{rec}_{\mathbb{B}^m} P P_t P_f \text{false}^m & \equiv P_f \\
\text{rec}_{\mathbb{B}^e} P P_t P_f \text{true}^e & \equiv P_t & \text{rec}_{\mathbb{B}^e} P P_t P_f \text{false}^e & \equiv P_f \\
\text{rec}_{\mathbb{B}^e} P P_t P_f (\text{raise } \mathbb{B}^e M) & \equiv \text{raise } (P(\text{raise } \mathbb{B}^e M)) M \\
\text{rec}_{\text{list}} A P P_n P_c (\text{nil } A) & \equiv P_n & \text{rec}_{\text{list}} A P P_n P_c (\text{cons } A M L) & \equiv P_c M L (\text{rec}_{\text{list}} A P P_n P_c L) \\
\text{rec}_{\text{eq}} A M P P_r M (\text{refl } A M) & \equiv P_r
\end{aligned}$$

Fig. 3. Examples of inductive types in various layers

Other than these specificities, the typing rules are standard. The rules for conversion, also standard, are given in Figure 2.

3.2 Inductive Types

Each layer e, m or p of RETT contains inductively generated types. Similarly to what happens in vanilla Coq with the Prop-Type distinction [Bertot and Castéran 2004], RETT inductive types and their elimination principles thus need to be specifically placed in a given “home” layer, $\square^e$, $\square^m$ or $\square^p$. When restricting the RETT system to a particular layer, this gives a full interpretation of CIC per hierarchy. We will thus not describe in detail their formation and elimination rules, as they are essentially the same as in CIC.

Inductive types meant to be used in programs must be placed in the exceptional or mediation layers. Conversely, inductive types meant for logical purposes must be placed in the mediation or pure layers. The dual role of the mediation layer will be explained thanks to the use of modalities later on. We give a few examples in Figure 3. To contrast the difference between the exceptional and mediation layers, we provide two variants of the booleans, $\mathbb{B}^e$ in $\square^e$ and $\mathbb{B}^m$ in $\square^m$. Note how the predicate of each eliminator lands in the same layer as the inductive type.

		Return type		
		$\Box^e$	$\Box^m$	$\Box^p$
Source	$\Box^e$	rec/catch	catch	catch
	$\Box^m$	rec	rec	–
	$\Box^p$	rec	rec	rec

Fig. 4. Legal mixed-layer eliminators

Note that because the exceptional layer features closed inductive terms that are not convertible to constructors, the reduction rules for eliminators over inductive types living in $\Box^e$ need to be extended to handle the `raise` term, by simply re-raising it. This is dictated by the usual semantics of call-by-name exceptions [Pédrot and Tabareau 2018].

On the other hand, the existence of additional exception-raising terms in $\Box^e$ is also reflected by the ability to handle exceptions on inductive types.

Catch Eliminators. Inductive types in the exceptional layer additionally feature a *catch eliminator*, which is the same as the standard eliminator extended with a premise for the `raise` term, which furthermore satisfy the expected equations of a `try/with` block.

For instance, effectful booleans are equipped with the constant

$$\text{catch}_{\mathbb{B}^e} : \Pi P : \mathbb{B}^e \rightarrow \Box_i^e. P \text{ true}^e \rightarrow P \text{ false}^e \rightarrow (\Pi e : E. P (\text{raise } \mathbb{B}^e e)) \rightarrow \Pi b : \mathbb{B}^e. P b$$

which is subject to the following equations

$$\begin{aligned} \text{catch}_{\mathbb{B}^e} P P_t P_f P_e \text{ true}^e &\equiv P_t & \text{catch}_{\mathbb{B}^e} P P_t P_f P_e \text{ false}^e &\equiv P_f \\ \text{catch}_{\mathbb{B}^e} P P_t P_f P_e (\text{raise } \mathbb{B}^e M) &\equiv P_e M \end{aligned}$$

Note that the usual eliminator for exceptional inductive types (e.g. $\text{rec}_{\mathbb{B}^e}$) can actually be derived from the catch eliminator by re-raising the exception in the handling branch. The resulting terms satisfy the expected equations automatically.

This generalized eliminator permits writing exception-handling code in the exceptional layer, as if this fragment was an impure programming language. If we were to use a pattern-matching based presentation, it would simply correspond to an optional additional exception-handling branch for exceptional inductive types.

Mixed Elimination. An even more interesting phenomenon at play is the interaction between the various layers. Indeed, eliminating an inductive type living in a layer s_1 into a layer s_2 needs to fulfill the invariants corresponding to their respective layers. This is not unlike what happens when eliminating from `Prop` to `Type` in CIC, insofar as one has to respect the singleton elimination criterion [Letouzey 2004]. Rather than having to decide that an elimination is proof-irrelevant, in RETT one has to check that an elimination cannot endanger consistency by allowing stray exceptions to land in a consistent layer.

This restriction is materialized by the fact that when eliminating from the unsafe $\Box^e$ layer to a safe layer, one has to explicitly handle all potential exceptions, by providing a catch-all clause. In practice, this means that there is no standard eliminator, only a catch eliminator. Given a layer for the eliminated inductive and a layer for the return predicate, the legal eliminators are summarized in Figure 4.

Note that for technical reasons that will be explained in the model construction (Section 4), it is not possible to eliminate from the mediation layer into the pure layer. Let us also insist that catch eliminators only make sense on exceptional inductive types, as the other layers lack a `raise` term, and the catch eliminator would not type-check.

3.3 Navigating Between Hierarchies

The main novelty of RETT is to provide modalities to navigate between the different layers, which are given below.

$$\{-\}_m^e : \Box^e \rightarrow \Box^m \quad \{-\}_e^m : \Box^m \rightarrow \Box^e \quad \{-\}_p^m : \Box^m \rightarrow \Box^p \quad \{-\}_m^p : \Box^p \rightarrow \Box^m$$

Although they are written in a uniform way, they have wildly different computational behaviors, reflecting the different properties of the three universe hierarchies.

The modality $\{-\}_m^e$ amounts to considering that all terms in an exceptional type A , including exceptions, are parametric when seen in $\{A\}_m^e$. Intuitively, one can understand these terms as suspended computations, which are therefore trivially harmless. The modality $\{-\}_e^m$ is just a forgetful functor, which forgets the notion of parametricity attached to a type in the mediation layer, thus releasing its ability to raise exceptions. The modality $\{-\}_p^m$ forgets both about parametricity and the ability to raise an exception at that type, seeing it as a pure type. The modality $\{-\}_m^p$ equips a pure type with a default way to raise an exception, but automatically forbidden by the equipped notion of parametricity.

Most notably, the two modalities originating from the mediation layer are well-behaved in the sense that they commute with type formers, while the other ones enjoy no such property. The reason is that behind the scenes, $\{-\}_e^m$ and $\{-\}_p^m$ are the only modalities which correspond to forgetful functors and do not add anything to the type. In particular, sending a mediation type into the exceptional layer results in an exceptional type.

$$\{\Box_i^m\}_e^m \equiv \Box_i^e$$

This equation should really be thought of as the fact that $\Box^e$ is a semantic supertype of $\Box^m$, or dually, that $\Box^m$ is a restriction of $\Box^e$. We insist that homologous conversions do not hold for any other modality. Likewise, the two modalities originating from the mediation layer commute with products.

$$\{\Pi x :^s A. B\}_e^m \equiv \Pi x :^s A. \{B\}_e^m \quad \{\Pi x :^s A. B\}_p^m \equiv \Pi x :^s A. \{B\}_p^m$$

These modalities are equipped with the corresponding introduction operators

$$\begin{aligned} \iota_m^e : \Pi A : \Box^e. A \rightarrow \{A\}_m^e & \quad \iota_e^m : \Pi A : \Box^m. A \rightarrow \{A\}_e^m \\ \iota_m^p : \Pi A : \Box^p. A \rightarrow \{A\}_m^p & \quad \iota_p^m : \Pi A : \Box^m. A \rightarrow \{A\}_p^m \end{aligned}$$

The corresponding equations hold on λ -abstractions. Note that the commutation of modalities with Π -types is necessary for these equations to be well-typed.

$$\begin{aligned} \iota_e^m (\Pi x :^s A. B) (\lambda x :^s A. M) & \equiv \lambda x :^s A. \iota_e^m B M \\ \iota_p^m (\Pi x :^s A. B) (\lambda x :^s A. M) & \equiv \lambda x :^s A. \iota_p^m B M \end{aligned}$$

The four modalities enjoy elimination principles, as long as the return type of the predicate is living in the same hierarchy as the target of the modality. Note that in general, there are no eliminators with a predicate living in a distinct layer.

$$\begin{aligned} elim_m^e & : \Pi(A : \Box^e) (P : \{A\}_m^e \rightarrow \Box^m). (\Pi a : A. P (\iota_m^e A a)) \rightarrow \Pi x : \{A\}_m^e. P x \\ elim_p^m & : \Pi(A : \Box^m) (P : \{A\}_p^m \rightarrow \Box^p). (\Pi a : A. P (\iota_p^m A a)) \rightarrow \Pi x : \{A\}_p^m. P x \\ elim_m^p & : \Pi(A : \Box^p) (P : \{A\}_m^p \rightarrow \Box^m). (\Pi a : A. P (\iota_m^p A a)) \rightarrow \Pi x : \{A\}_m^p. P x \\ elim_e^m & : \Pi(A : \Box^m) (P : \{A\}_e^m \rightarrow \Box^e). (\Pi a : A. P (\iota_e^m A a)) \rightarrow \Pi x : \{A\}_e^m. P x \end{aligned}$$

These eliminators satisfy the expected reduction rules, e.g.

$$elim_m^e A P P_i (\iota_m^e A M) \equiv P_i M$$

Going from $\Box^e$ to $\Box^m$ and going back is the identity because it equips an exceptional type with a default notion of parametricity and directly forgets it, which is formally described by the following conversions:

$$\{\{A\}_m^e\}_e^m \equiv A \quad \text{and} \quad \iota_e^m \{A\}_m^e (\iota_m^e A M) \equiv M.$$

3.4 Internal Parametricity

RETT also features an internal parametricity predicate $\mathcal{P}$ on types of the form $\{A\}_e^m$, which classifies effectful terms that happen to be pure. It comes equipped with an injection and a projection

$$\begin{aligned} \mathcal{P} & : \Pi A : \Box^m. \{A\}_e^m \rightarrow \Box^m \\ \iota_{\mathcal{P}} & : \Pi(A : \Box^m) (a : A). \mathcal{P} A (\iota_e^m A a) \\ \Downarrow_{\mathcal{P}} & : \Pi(A : \Box^m) (a : \{A\}_e^m). \mathcal{P} A a \rightarrow A \end{aligned}$$

that satisfy the following equations

$$\Downarrow_{\mathcal{P}} A (\iota_e^m A M) (\iota_{\mathcal{P}} A M) \equiv M \quad \iota_e^m A (\Downarrow_{\mathcal{P}} A M P) \equiv M$$

Interestingly, using these terms, we can provide a specific elimination principle $\text{elim}_{\mathcal{P}}$ that enables reasoning on $\{-\}_e^m$ with predicates living in the mediation layer. As such, it is the mediation-landing version of the elim_e^m eliminator, and intuitively the requirement that the term being eliminated is parametric corresponds to the invariant that it should not raise uncaught exceptions.

The parametricity eliminator is defined as

$$\begin{aligned} \text{elim}_{\mathcal{P}} & : \Pi(A : \Box^m) (P : \{A\}_e^m \rightarrow \Box^m). (\Pi x : A. P (\iota_e^m A x)) \rightarrow \Pi(x : \{A\}_e^m) (p : \mathcal{P} A x). P x \\ & := \lambda A P P_i x p. P_i (\Downarrow_{\mathcal{P}} A x p). \end{aligned}$$

LEMMA 3.1. *The parametricity eliminator satisfies the expected ι -reduction:*

$$\text{elim}_{\mathcal{P}} A P P_i (\iota_e^m A M) (\iota_{\mathcal{P}} A M) \equiv P_i M.$$

Internal parametricity lets the user separate the implementation of a term that potentially uses exceptions internally from its specification, ensuring that it is observationally pure. We use it in the next section to derive standard elimination of exceptional inductive types into the mediation layer, up to a proof of parametricity of the exceptional term.

Similarly to what happens with corresponding modalities, the parametricity predicate lives in the mediation layer, and thus commutes with dependent products.

LEMMA 3.2. *Using the above combinators, it is possible to write terms*

$$\begin{aligned} \mathcal{P}_{\text{to_}\Pi} & : \Pi(A : \Box^s) (B : A \rightarrow \Box^m) (f : \{\Pi x :^s A. B x\}_e^m). \mathcal{P} (\Pi x :^s A. B x) f \rightarrow \Pi x :^s A. \mathcal{P} (B x) (f x) \\ \mathcal{P}_{\text{of_}\Pi} & : \Pi(A : \Box^s) (B : A \rightarrow \Box^m) (f : \{\Pi x :^s A. B x\}_e^m). (\Pi x :^s A. \mathcal{P} (B x) (f x)) \rightarrow \mathcal{P} (\Pi x :^s A. B x) f \end{aligned}$$

Those terms satisfy unsurprising equations (not detailed here) that can be used to easily transfer parametricity conditions under and over contexts.

3.5 Parametricity for Exceptional Inductive Types

We now show that the $\{-\}_e^m$ modality together with the parametricity predicate $\mathcal{P}$ make it possible to define a notion of parametricity on exceptional inductive types. The notion of parametricity in turn allows us to derive standard elimination principles *into the mediation layer* for exceptional inductive types, with an extra guard condition that the eliminated exceptional term is parametric.

Definition 3.3. Let $\mathcal{I} : \Pi(x_1 :^{s_1} X_1) \dots (x_n :^{s_n} X_n). \Box^m$ be an inductive type in the mediation layer. We call $\{\mathcal{I}\}_e^m$ the *exceptional lowering* of $\mathcal{I}$, which is typed as $\{\mathcal{I}\}_e^m : \Pi(x_1 :^{s_1} X_1) \dots (x_n :^{s_n} X_n). \Box^e$.

We now show how the lowering of an inductive type from the mediation layer (called a mediation inductive type for short) satisfies the parametric elimination principle mentioned above and is equivalent to its corresponding exceptional inductive type in the case of booleans and lists.

Non-recursive types. Lowering a non-recursive mediation inductive type through $\{-\}_e^m$ results in an inductive type (e.g. $\{\mathbb{B}^m\}_e^m$) that behaves just like an exceptional inductive type (e.g. $\mathbb{B}^e$). Lowered inductive types can be introduced by the corresponding injection of their constructors. For instance, in the case of booleans, we have

$$\iota_e^m \mathbb{B}^m \text{true}^m : \{\mathbb{B}^m\}_e^m \quad \text{and} \quad \iota_e^m \mathbb{B}^m \text{false}^m : \{\mathbb{B}^m\}_e^m.$$

Usual eliminators targetting the exceptional layer can be derived using the eliminators for the corresponding mediation inductive together with the eliminator for the lowering modality. We can e.g. implement a term

$$\text{rec}_{\{\mathbb{B}^m\}_e^m} : \Pi P : \{\mathbb{B}^m\}_e^m \rightarrow \square^e. P (\iota_e^m \mathbb{B}^m \text{true}^m) \rightarrow P (\iota_e^m \mathbb{B}^m \text{false}^m) \rightarrow \Pi b : \{\mathbb{B}^m\}_e^m. P b$$

satisfying the expected ι -rules for constructors *only*. The reduction rule of the modality eliminator on raise is indeed not specified, which prevents extending the equation on raise to the lowered inductive version.

By restricting oneself to the case of *parametric* inductive terms, it is also possible to write an eliminator that targets the mediation layer. It is readily implemented by chaining the eliminator for internal parametricity with the one for the inductive type under consideration. For booleans, this results in a *parametric eliminator*

$$\text{rec}_{\{\mathbb{B}^m\}_e^m}^{\mathcal{P}} : \Pi P : \{\mathbb{B}^m\}_e^m \rightarrow \square^m. P (\iota_e^m \mathbb{B}^m \text{true}^m) \rightarrow P (\iota_e^m \mathbb{B}^m \text{false}^m) \rightarrow \Pi b : \{\mathbb{B}^m\}_e^m. \mathcal{P} \mathbb{B}^m b \rightarrow P b$$

that is also subject to the expected conversion rules. Such parametric eliminators allow us to reason in the mediation layer about the purity of terms of exceptional inductive types.

Unfortunately, catch eliminators for lowered inductive types are not derivable from the set of primitive combinators at hand. Thankfully, lowered catch eliminators are nonetheless valid in the model (i.e. one can provide a term in the target theory whose type is the translation of the corresponding source type), and thus can be postulated in RETT. For booleans, this amounts to stating that RETT is extended with terms of type

$$\begin{aligned} \text{catch}_{\{\mathbb{B}^m\}_e^m} : \Pi P : \{\mathbb{B}^m\}_e^m \rightarrow \square^s. \\ P (\iota_e^m \mathbb{B}^m \text{true}^m) \rightarrow P (\iota_e^m \mathbb{B}^m \text{false}^m) \rightarrow (\Pi e : E. P (\text{raise } \{\mathbb{B}^m\}_e^m e)) \rightarrow \\ \Pi b : \{\mathbb{B}^m\}_e^m. P b \end{aligned}$$

for s ranging over e, m and p and subject to the full range of equations, that is, both the ones on constructors as well as the ones on raise .

One can now use these lowered catch eliminators to show that the lowering of a mediation inductive type is isomorphic to the corresponding exceptional inductive type. This makes explicit the dual nature of the mediation layer, which can be used both for safe reasoning, and for effectful computation through lowering.

LEMMA 3.4. $\{\mathbb{B}^m\}_e^m$ is isomorphic to $\mathbb{B}^e$.

PROOF. The two inductive types satisfy the same universal property, namely catch elimination, and thus are isomorphic. $\square$

Moreover, the ability to catch failures on lowered inductive types can also be used to specify the parametricity predicate on them. That is, it is possible to prove that failure on lowered inductive types is never parametric¹, e.g. for booleans

LEMMA 3.5. *The following type is inhabited in RETT*

$$\Pi (P : \{\mathbb{B}^m\}_e^m \rightarrow \square^m) (e : E). \mathcal{P} \mathbb{B}^m (\text{raise } \{\mathbb{B}^m\}_e^m e) \rightarrow P (\text{raise } \{\mathbb{B}^m\}_e^m e).$$

¹This is equivalent to saying that $\mathcal{P} \mathbb{B}^m (\text{raise } \{\mathbb{B}^m\}_e^m e)$ implies $\perp_m$, the inductive with no constructor in $\square^m$.

This is easily obtained by combining the parametric eliminator with the catch eliminator.

Recursive types. We conclude this section with the specific case of lowering recursive inductive types, i.e. types that mention themselves in the type of their constructors. In this case, one has to be a little more careful than above, because lowering needs to be handled specially. The reason is that the $\{-\}_e^m$ modality does not distribute on the left-hand side of a Π -type, which means that there is a type mismatch for recursive constructors, e.g.

$$\iota_e^m (A \rightarrow \text{list } A \rightarrow \text{list } A) \text{ cons} : A \rightarrow \text{list } A \rightarrow \{\text{list } A\}_e^m$$

rather than

$$A \rightarrow \{\text{list } A\}_e^m \rightarrow \{\text{list } A\}_e^m$$

which would be required to obtain an inductive equivalent to the list datatype in $\square^e$. This can be circumvented using the elimination principle of $\{-\}_e^m$ on the recursive arguments. For instance,

$$\text{elim}_e^m (\text{list } A) (\lambda _ . \{\text{list } A\}_e^m) (\lambda l . \iota_e^m (\text{list } A) (\text{cons } A \ M \ l)) \ N$$

has the adequate type as expected above.

4 A SYNTACTIC MODEL OF RETT

We define the semantics of RETT by a syntactic translation into CIC, following the general technique of [Boulier et al. \[2017\]](#). This allows us to straightforwardly prove its good metatheoretical properties, like consistency and canonicity. We first present the translations of each layer, then explain the translation of the modalities and finally prove the correctness of the translation and deduce metatheoretical properties.

4.1 Exceptional Layer

The translation of the exceptional layer is given in Figure 5. It follows exactly the translation given by [Pédrot and Tabareau \[2018\]](#), which we almost completely introduced in Section 2. Following the syntactic translation approach [\[Boulier et al. 2017\]](#), the term translation is written $[-]^e$ and the type translation, written $\llbracket - \rrbracket^e$, is derived from it using the function El .² Note that in RETT the exceptional translation only applies to terms whose type is a sort in the exceptional universe $\square^e$.

Recall that types are mapped to values of the inductive type type, which has two constructors, `TypeVal` and `TypeErr`. The former is used to represent valid types (as a pair of a type and its default function); the latter is the default function for errors on types. The only rule we did not explain in Section 2 is the translation of the dependent product: it simply produces a `TypeVal` whose representation type is the type component of the recursive translation on A and B , and whose default function re-raises the exception e on the default function for type B (retrieved using the macro $[-]_\emptyset$).

The translation handles inductive types following the approach of [Pédrot and Tabareau \[2018\]](#) briefly presented in Section 2. We provide the examples of booleans and lists for illustration. Essentially, an inductive type (e.g. $\mathbb{B}^e$) is translated to a new inductive type (e.g. $\mathbb{B}^\bullet$) with an extra constructor (e.g. $\mathbb{B}_\emptyset$), used as the default function to raise exceptions at that type. The eliminators (e.g. $\text{rec}_\mathbb{B}$) propagate exceptions in these new branches.

4.2 Mediation Layer

The translation $[-]^m$ of the mediation layer is the same as that of the exceptional layer, replacing e with m , in particular:

²Recall from Section 2 that El_i recovers the underlying type from an inhabitant of type_i , and Err_i lifts the default function to this underlying type.

$$\begin{aligned}
[\Box_i^e]^e &:= \text{TypeVal } \text{type}_i \text{ TypeErr}_i \\
[x]^e &:= x \\
[\lambda x :^s A. M]^e &:= \lambda x : \llbracket A \rrbracket^s. [M]^e \\
[M^s N]^e &:= [M]^e [N]^s \\
[\Pi x :^s A. B]^e &:= \text{TypeVal } (\Pi x : \llbracket A \rrbracket^s. \llbracket B \rrbracket^e) (\lambda(e : \mathbb{E}) (x : \llbracket A \rrbracket^s). [B]_{\emptyset} e) \\
[\mathbb{E}]^e &:= \text{TypeVal } \mathbb{E} (\lambda e : \mathbb{E}. e) \\
[\text{raise}]^e &:= \lambda(A : \text{type}) (e : \mathbb{E}). [A]_{\emptyset} e \\
[\mathbb{B}]^e &:= \text{TypeVal } \mathbb{B}^{\bullet} \mathbb{B}_{\emptyset} \\
[\text{list}]^e &:= \lambda A : \llbracket \Box^e \rrbracket. \text{TypeVal } (\text{list}^{\bullet} [A]) \text{list}_{\emptyset} \\
[c]^e &:= c^{\bullet} \quad (\text{for any constructor } c \text{ of an inductive type}) \\
[\text{rec}_{\mathbb{B}}]^e &:= \lambda P p_t p_f b. \text{ match } b \text{ return } \lambda b. \text{El } (P b) \text{ with} \\
&\quad | \text{true}^{\bullet} \Rightarrow p_t \\
&\quad | \text{false}^{\bullet} \Rightarrow p_f \\
&\quad | \mathbb{B}_{\emptyset} e \Rightarrow [P \mathbb{B}_{\emptyset} e]_{\emptyset} e \\
&\quad \text{end} \\
[\text{rec}_{\text{list}}]^e &:= \dots \text{ (omitted for brevity)} \\
[A]_{\emptyset} &:= \text{Err } [A]^e \\
\llbracket A \rrbracket^e &:= \text{El } [A]^e \\
\llbracket \cdot \rrbracket &:= \cdot \\
\llbracket \Gamma, x :^e A \rrbracket &:= \llbracket \Gamma \rrbracket, x : \llbracket A \rrbracket^e
\end{aligned}$$

$$\begin{array}{l|l}
\text{Ind } \mathbb{B}^{\bullet} : \Box := & \text{Ind list}^{\bullet} (A : \llbracket \Box^e \rrbracket) : \Box := \\
| \text{true}^{\bullet} : \mathbb{B}^{\bullet} & | \text{nil}^{\bullet} : \text{list}^{\bullet} A \\
| \text{false}^{\bullet} : \mathbb{B}^{\bullet} & | \text{cons}^{\bullet} : \llbracket A \rrbracket \rightarrow \text{list}^{\bullet} A \rightarrow \text{list}^{\bullet} A \\
| \mathbb{B}_{\emptyset} : \mathbb{E} \rightarrow \mathbb{B}^{\bullet} & | \text{list}_{\emptyset} : \mathbb{E} \rightarrow \text{list}^{\bullet} A
\end{array}$$

Fig. 5. Translation of the exceptional layer

$$\begin{aligned}
[\Box_i^m]^m &:= \text{TypeVal } \text{type}_i \text{ TypeErr}_i \\
\llbracket A \rrbracket^m &:= \text{El } [A]^m
\end{aligned}$$

The peculiarity of the mediation layer is that every term also comes with its *parametricity proof*. This proof is obtained by the translation $[-]_{\varepsilon}$, described in Figure 6. This translation is essentially the standard parametricity translation for type theory [Bernardy and Lasson 2011], with a few adjustments specific to RETT. We stress that $[-]_{\varepsilon}$ is only defined for terms living in the mediation layer, so that writing $[M]_{\varepsilon}$ implicitly assumes $M : A$ for some $A : \Box^m$.

Let us first recall the basics of this translation. For the universe, the translation is defined as (arbitrary) predicates on types, i.e. if $A : \Box^m$ then $[A]_{\varepsilon} : \llbracket A \rrbracket^m \rightarrow \Box$. Dependent products, functions,

$$\begin{aligned}
[\Box_i^m]_\epsilon &:= \lambda A : \llbracket \Box_i^m \rrbracket. \llbracket A \rrbracket^m \rightarrow \Box_i \\
[x]_\epsilon &:= x_\epsilon \\
[\lambda x :^m A. M]_\epsilon &:= \lambda (x : \llbracket A \rrbracket^m) (x_\epsilon : \llbracket A \rrbracket_\epsilon x). [M]_\epsilon \\
[\lambda x :^s A. M]_\epsilon &:= \lambda x : \llbracket A \rrbracket^s. [M]_\epsilon \quad \text{if } s \in \{p, e\} \\
[M^m N]_\epsilon &:= [M]_\epsilon [N]^m [N]_\epsilon \\
[M^s N]_\epsilon &:= [M]_\epsilon [N]^s \quad \text{if } s \in \{p, e\} \\
[\Pi x :^m A. B]_\epsilon &:= \lambda (f : \Pi x : \llbracket A \rrbracket^m. \llbracket B \rrbracket^m). \Pi (x : \llbracket A \rrbracket^m) (x_\epsilon : \llbracket A \rrbracket_\epsilon x). \llbracket B \rrbracket_\epsilon (f x) \\
[\Pi x :^s A. B]_\epsilon &:= \lambda (f : \Pi x : \llbracket A \rrbracket^s. \llbracket B \rrbracket^m). \Pi x : \llbracket A \rrbracket^s. \llbracket B \rrbracket_\epsilon (f x) \quad \text{if } s \in \{p, e\} \\
[\mathcal{I}]_\epsilon &:= \mathcal{I}_\epsilon \quad (\text{for any inductive type } \mathcal{I}) \\
[c]_\epsilon &:= c_\epsilon \quad (\text{for any constructor } c \text{ of an inductive type}) \\
\llbracket A \rrbracket_\epsilon &:= [A]_\epsilon \\
\llbracket \cdot \rrbracket_\epsilon &:= \cdot \\
\llbracket \Gamma, x :^m A \rrbracket_\epsilon &:= \llbracket \Gamma \rrbracket_\epsilon, x : \llbracket A \rrbracket^m, x_\epsilon : \llbracket A \rrbracket_\epsilon x \\
\llbracket \Gamma, x :^s A \rrbracket_\epsilon &:= \llbracket \Gamma \rrbracket_\epsilon, x : \llbracket A \rrbracket^s \quad \text{if } s \in \{p, e\}
\end{aligned}$$

$$\begin{array}{l}
\text{Ind } \mathbb{B}_\epsilon : \mathbb{B}^\bullet \rightarrow \Box := \\
| \text{true}_\epsilon : \mathbb{B}_\epsilon \text{true}^\bullet \\
| \text{false}_\epsilon : \mathbb{B}_\epsilon \text{false}^\bullet
\end{array}
\quad \left\| \quad
\begin{array}{l}
\text{Ind } \text{list}_\epsilon (A : \text{type}) (A_\epsilon : \llbracket A \rrbracket \rightarrow \Box) : \text{list}^\bullet A \rightarrow \Box := \\
| \text{nil}_\epsilon : \text{list}_\epsilon A A_\epsilon (\text{nil}^\bullet A) \\
| \text{cons}_\epsilon : \Pi (x : \llbracket A \rrbracket) (x_\epsilon : A_\epsilon x) (l : \text{list}^\bullet A) (l_\epsilon : \text{list}_\epsilon A A_\epsilon l). \\
\quad \text{list}_\epsilon A A_\epsilon (\text{cons}^\bullet A x l)
\end{array}
\right.$$

Fig. 6. Parametricity translation

and applications are defined by cases, but consider only the first line of each for now. For the dependent product, the translation specifies that given a parametric input x of type A —as witnessed by x_ϵ of type $\llbracket A \rrbracket_\epsilon x$ —the function yields a parametric output of type B . Similarly, the translation of a lambda term is a function that takes an argument x and a witness x_ϵ that it is parametric; a variable x is translated to x_ϵ ; a translated application (again, consider only the first line for now) passes the parametricity witness as an extra argument. The translation of type environments follows the same pattern, with parametricity witness x_ϵ .

The main specificity of the parametricity translation for RETT is that it must take into account the fact the dependent product in RETT can quantify over types in any layer, in particular those which are not coming with parametricity proofs. Therefore, the translation of the dependent product depends on the layer of the domain: if $x : A$ is in the mediation layer, then the parametricity predicate for its argument (x_ϵ) is required; otherwise, it is not. The translations of lambda terms and applications follow the same discipline: parametricity is only imposed on types and terms from the mediation layer. Second, as in ETT, the parametricity translation recursively triggers the base translation $[-]^s$ (or $\llbracket - \rrbracket^s$ depending on the position) on any occurrence of a RETT term from the layer s —see for instance the translation of an application, which uses $[N]^s$.

The parametric translation of inductive types, illustrated for booleans and lists, differs from the exceptional in two crucial ways: first, no default constructors are added. This is because the

$$\begin{aligned}
[\Box_i^p]^p &:= \Box_i \\
[x]^p &:= x \\
[\lambda x :^s A. M]^p &:= \lambda x : \llbracket A \rrbracket^s. [M]^p \\
[M^s N]^p &:= [M]^p [N]^s \\
[\Pi x :^s A. B]^p &:= \Pi x : \llbracket A \rrbracket^s. [B]^p \\
[\Sigma x : A. B]^p &:= \Sigma x : \llbracket A \rrbracket^p. [B]^p \\
[\text{eq } A \ x \ y]^p &:= \text{eq } \llbracket A \rrbracket^p [x]^p [y]^p \\
\llbracket A \rrbracket^p &:= [A]^p
\end{aligned}$$

Fig. 7. Translation of the pure layer

parametric translation imposes purity, and hence only the standard constructors are valid, assuming their arguments are. This latter condition means that parametricity witnesses are required: e.g., cons_ε requires the parametricity witness of both the added element (x_ε) and the list (l_ε).

4.3 Pure Layer

The pure translation (Fig. 7) is essentially the identity translation, but for the fact that it inductively makes use of previous translation when using a crosscutting dependent product whose domain is in another layer.

4.4 Mixed Eliminators

The interpretation of inductive types in each layer is given above, but it is worth mentioning the crosscutting eliminators. Indeed, as discussed in Section 3.2, there are ways to eliminate inductive terms on predicates landing in a different layer than the one the inductive type is living in.

The various catch eliminators are simply built out of the corresponding pattern-matching on all constructors. For instance, the $\mathbb{B}^e$ eliminator into $\Box^m$ is defined in Figure 8. Note the intricate return type of the parametric component. Likewise, the catch eliminator into $\Box^e$ has the same base translation, the main difference being that it does not have a $[-]_\varepsilon$ translation. The catch eliminator into $\Box^p$ is also very similar, the main difference being the removal of E1 casts.

We will not describe the other eliminators as they are straightforward. We will insist nonetheless on the reason why there is no eliminator from $\Box^m$ inductive types into $\Box^p$. This is due to the lack of internal parametricity in $\Box^p$. Given a value of $\mathbb{B}^m$, through the $[-]^m$ one also needs to handle the failure case in the $\Box^p$ -returning pattern-matching, but there is no way to return a default value because in general types living in $\Box^p$ are not necessarily inhabited. One could then argue that it would still be possible to provide the catch variant. While it seems reasonable from a computational point of view, the problem is now that there is no way to give a RETT type to the failure premise, as the term $\text{raise } \mathbb{B}^m \ M$ is ill-typed. As a consequence, there is no catch term from $\Box^m$ into $\Box^p$.

4.5 Modalities

The translation of modalities is given in Figure 9. Notice that the composed modality that goes from $\Box^e$ to $\Box^p$ and back is not at all the identity, as it will add freely exceptions to a type that is already exceptional. This justifies the claim in Section 3 that no reasonable interplay is possible between the pure layer and the exceptional one. Indeed, adding exceptions to CIC is a whole

$$\begin{aligned}
[\text{catch}_{\mathbb{B}^e}] &:= \lambda P : \mathbb{B}^e \rightarrow \text{type}. \\
&\quad \lambda(P_t : \text{El } (P \text{ true}^\bullet)) (P_f : \text{El } (P \text{ false}^\bullet)) (P_e : \Pi e : \mathbb{E}. \text{El } (P (\mathbb{B}_\emptyset e))). \\
&\quad \lambda b : \mathbb{B}^e. \\
&\quad \text{match } b \text{ return } \lambda b. \text{El } (P b) \text{ with} \\
&\quad \quad | \text{true}^\bullet \Rightarrow P_t \\
&\quad \quad | \text{false}^\bullet \Rightarrow P_f \\
&\quad \quad | \mathbb{B}_\emptyset e \Rightarrow P_e e \\
&\quad \text{end} \\
[\text{catch}_{\mathbb{B}^e}]_\epsilon &:= \lambda(P : \mathbb{B}^e \rightarrow \text{type}) (P_\epsilon : \Pi b : \mathbb{B}^e. \text{El } (P b) \rightarrow \square). \\
&\quad \lambda(P_t : \text{El } (P \text{ true}^\bullet)) (P_{t_\epsilon} : P_\epsilon \text{ true}^\bullet P_t). \\
&\quad \lambda(P_f : \text{El } (P \text{ false}^\bullet)) (P_{f_\epsilon} : P_\epsilon \text{ false}^\bullet P_f). \\
&\quad \lambda(P_e : \Pi e : \mathbb{E}. \text{El } (P (\mathbb{B}_\emptyset e))) (P_{e_\epsilon} : \Pi e : \mathbb{E}. P_\epsilon (\mathbb{B}_\emptyset e) (P_e e)). \\
&\quad \lambda b : \mathbb{B}^e. \\
&\quad \text{match } b \text{ return } \lambda b. P_\epsilon b ([\text{catch}_{\mathbb{B}^e}] P P_t P_f P_e b) \text{ with} \\
&\quad \quad | \text{true}^\bullet \Rightarrow P_{t_\epsilon} \\
&\quad \quad | \text{false}^\bullet \Rightarrow P_{f_\epsilon} \\
&\quad \quad | \mathbb{B}_\emptyset e \Rightarrow P_{e_\epsilon} e (P_e e) \\
&\quad \text{end}
\end{aligned}$$
Fig. 8. Eliminating $\mathbb{B}^e$ into $\square^m$

program translation that deeply modifies the structure of the program so that one cannot *internally* go back and forth between the two layers while preserving the program structure.

The term translation of $\{-\}_m^m$ is the identity, as it only consists in forgetting the parametricity witnesses when going from $\square^m$ to $\square^e$. The translation of $\{-\}_p^m$ is given by El as it consists in recovering the underlying type of an inhabitant of type.³ The translation of $\{-\}_m^p$ is given by freely adding the exception type $\mathbb{E}$ to the base pure type using a sum type. Its parametricity predicate corresponds to witnesses that the inhabitant of the sum type is actually a value rather than an exception. To this end, we use the following dedicated inductive types in the target theory.

Inductive $\mathcal{E} (A : \square) : \square := \text{val} : A \rightarrow \mathcal{E} A \mid \text{err} : \mathbb{E} \rightarrow \mathcal{E} A$
Inductive $\text{IsV} (A : \square) : \mathcal{E} A \rightarrow \square := \text{isV} : \Pi a : A. \text{IsV } A (\text{val } A a)$

Finally, the translation of $\{-\}_m^e$ is given by the identity on the $[-]$ part; its parametricity predicate is trivial, as described by the following inductive type in the target theory.

Inductive $\top : \square := \text{tt} : \top$

The translation of the introduction operators, presented in Figure 10, is straightforward, being either the identity or a canonical injection. Note in particular that $[l_m^e A (\text{raise } A e)]_\epsilon \equiv \text{tt}$: exceptions can live in the mediating layer through the modality, as trivially harmless terms.

The translation of eliminators (Fig. 11) is more complex. The translations of elim_e^m and elim_p^m are almost the identity. The translations of elim_m^e and elim_m^e require the use of the eliminators to the inductive types introduced by the translation. $[\text{elim}_m^p]$ pattern-matches on the inhabitant of x of type $\mathcal{E} [A]$: if it is a value (i.e. of the form $\text{val } A a$), it uses P_a given in the hypothesis; if it

³We would like to define the translation as the combination of an element of the underlying type plus a proof that it is parametric, but we do not have access to the parametricity predicate in the $[-]$ translation. This definition will be made possible in Section 5 by considering a subtheory of RETT.

$$\begin{array}{l|l}
[\{A\}_e^m]^e & := [A]^m \\
[\{A\}_p^m]^p & := \text{El } [A]^m \\
[\{A\}_m^p]^m & := \text{TypeVal } (\mathcal{E} \llbracket A \rrbracket^p) (\text{err } \llbracket A \rrbracket^p) \\
[\{A\}_m^e]^e & := [A]^e
\end{array}
\quad
\begin{array}{l}
[\{A\}_m^p]_\varepsilon^p := \lambda x : \mathcal{E} \llbracket A \rrbracket^p. \text{IsV } \llbracket A \rrbracket^p x \\
[\{A\}_m^e]_\varepsilon^e := \lambda x : \text{El } [A]^e. \top
\end{array}$$

Fig. 9. Translation of the four modalities

$$\begin{array}{l|l}
[\iota_e^m]^e & := \lambda(A : \text{type}) (x : \text{El } A). x \\
[\iota_p^m]^p & := \lambda(A : \text{type}) (x : \text{El } A). x \\
[\iota_m^p]^m & := \lambda(A : \square) (x : A). \text{val } A x \\
[\iota_m^e]^e & := \lambda(A : \text{type}) (x : \text{El } A). x
\end{array}
\quad
\begin{array}{l}
[\iota_m^p]_\varepsilon^p := \lambda(A : \square) (a : A). \text{isV } A a \\
[\iota_m^e]_\varepsilon^e := \lambda(A : \text{type}) (_ : \text{El } A). \text{tt}
\end{array}$$

Fig. 10. Translation of the four introduction operators

$$\begin{array}{l}
[\text{elim}_e^m]^e := \lambda(A : \text{type}) (P : \text{El } A \rightarrow \text{type}) (P_a : \Pi a : \text{El } A. \text{El } (P a)) (x : \text{El } A). P_a x \\
[\text{elim}_p^m]^p := \lambda(A : \text{type}) (P : \text{El } A \rightarrow \square) (P_a : \Pi a : \text{El } A. P a) (x : \text{El } A). P_a x \\
[\text{elim}_m^p]^m := \lambda(A : \square) (P : \mathcal{E} A \rightarrow \text{type}) (P_a : \Pi a : A. \text{El } (P (\text{val } A a))) (x : \mathcal{E} A). \\
\quad \text{rec}_{\mathcal{E}} P (\lambda a. P_a a) (\lambda e. \text{Err } (P (\text{err } e))) x \\
[\text{elim}_m^e]^m := \lambda(A : \text{type}) (P : \text{El } A \rightarrow \text{type}) (P_a : \Pi a : \text{El } A. \text{El } (P a)) (x : \text{El } A). P_a x \\
[\text{elim}_m^p]_\varepsilon^p := \lambda(A : \square) (P : \mathcal{E} A \rightarrow \text{type}) (P_\varepsilon : \Pi(x : \mathcal{E} A) x_\varepsilon. \text{El } (P x) \rightarrow \square). \\
\quad \lambda(P_a : \Pi a : A. \text{El } (P (\text{val } a))) (P_{a_\varepsilon} : \Pi a : A. P_\varepsilon (\text{val } A a) (\text{isV } A a) (P_a a)). \\
\quad \lambda(x : \mathcal{E} A) (x_\varepsilon : \text{IsV } A x). \\
\quad \text{rec}_{\text{IsV } A} (\lambda(x : \mathcal{E} A) x_\varepsilon. P_\varepsilon x x_\varepsilon ([\text{elim}_m^p]^m A P P_a x)) P_{a_\varepsilon} x x_\varepsilon \\
[\text{elim}_m^e]_\varepsilon^e := \lambda(A : \text{type}) (P : \text{El } A \rightarrow \text{type}) (P_\varepsilon : \Pi(a : \text{El } A) a_\varepsilon. \text{El } (P a) \rightarrow \square). \\
\quad \lambda(P_a : \Pi a : \text{El } A. \text{El } (P a)) (P_{a_\varepsilon} : \Pi a : \text{El } A. P_\varepsilon a \text{ tt } (P_a a)) (x : \text{El } A) (x_\varepsilon : \top). \\
\quad \text{rec}_\top (\lambda p. P_\varepsilon x p (P_a x)) (P_{a_\varepsilon} x) x_\varepsilon
\end{array}$$

Fig. 11. Translation of the four eliminators of modalities

is an exception (i.e. of the form $\text{err } A e$) it re-raises the exception of the return predicate. $[\text{elim}_m^e]$ is almost the identity. $[\text{elim}_m^p]_\varepsilon^p$ pattern-matches on the parametricity proof, which ensures that a parametric inhabitant of $\mathcal{E} [A]$ is equal to a term $\text{val } A a$ for some a in A . The translation of $[\text{elim}_m^e]_\varepsilon^e$ is given by the fact that any x_ε in $\top$ is equal to tt .

Note that the translation of modalities allows us to show a variant of Lemma 3.4, that the translation of an exceptional inductive type and its corresponding lowering are convertible.

LEMMA 4.1. $[\{\mathbb{B}^m\}_e^m]^e \equiv [\mathbb{B}^e]^e$.

4.6 Parametricity Predicate

As explained in Section 3.3, any type in the exceptional layer of the form $\{A\}_e^m$ can be equipped with a parametricity predicate $\mathcal{P} A$ coming from $[A]_\varepsilon$. The translation of $[\mathcal{P} A]$ is just given by $\top$ as there

$$\begin{aligned}
[\mathcal{P}]^m &:= \lambda(A : \text{type}) (x : \text{El } A). \top \\
[\iota_{\mathcal{P}}]^e &:= \lambda(A : \text{type}) (x : \text{El } A). \text{tt} \\
[\Downarrow_{\mathcal{P}}]^m &:= \lambda(A : \text{type}) (x : \text{El } A) (p : \top). x \\
[\mathcal{P}]_{\varepsilon} &:= \lambda(A : \text{type}) (A_{\varepsilon} : \text{El } A \rightarrow \square) (x : \text{El } A). \top \rightarrow A_{\varepsilon} x \\
[\Downarrow_{\mathcal{P}}]_{\varepsilon} &:= \lambda(A : \text{type}) (A_{\varepsilon} : \text{El } A \rightarrow \square) (x : \text{El } A) (p : \top) (p_{\varepsilon} : A_{\varepsilon} x). p_{\varepsilon}
\end{aligned}$$

Fig. 12. Translation of the $\mathcal{P}$ predicate

is no information to provide at this stage, the parametricity predicate being available only for the parametric translation. The translation of $[\mathcal{P} A]_{\varepsilon}$ is simply returning the parametricity predicate $[A]_{\varepsilon}$ given by the translation. The translation of the elimination principle of $\mathcal{P}$ is straightforward. The parametricity predicate of types lifted from the exceptional layer is trivial, as both $[\mathcal{P} \{A\}_{\varepsilon}^m] \equiv \top$ and $[\mathcal{P} \{A\}_{\varepsilon}^m]_{\varepsilon} \equiv \top$, which explains why there is no way to extract any useful content from it.

Finally, any inductive type in the mediation layer (e.g. $\mathbb{B}^m$) gives rise to a catch recursor on its lowering (e.g. $\{\mathbb{B}^m\}_{\varepsilon}^m$). Due to the fact that there is no difference in the underlying translation between $\mathbb{B}^m$ and $\mathbb{B}^e$, the translation of this recursor is the same as in Section 4.4, that is, it is given by pattern-matching with the additional failure case being handled by the additional failure premise.

4.7 Metatheoretical Properties of RETT

The soundness of the translations $[-]^s$ follow from the following properties.

THEOREM 4.2 (SOUNDNESS). *The following properties hold.*

- $[M\{x := N\}]^s \equiv [M]^s \{x := [N]^s\}$ (substitution lemma).
- *If $M \equiv N$ then $[M]^s \equiv [N]^s$ (conversion lemma).*
- *If $\Gamma \vdash M : A$ then $\llbracket \Gamma \rrbracket \vdash [M]^s : \llbracket A \rrbracket^s$ (typing soundness).*
- *If $\Gamma \vdash A : \square^s$ then $\llbracket \Gamma \rrbracket \vdash [A]_{\emptyset} : \mathbb{E} \rightarrow \llbracket A \rrbracket^s$, when $s \in \{e, m\}$ (exception soundness).*

PROOF. The first property is by routine induction on M , the second is direct by induction on the conversion derivation. The third is by induction on the typing derivation. As in [Pédrot and Tabareau 2018], the most important rule is $\square_i^s : \square_j^s$, for the three layers. For the exception layer (and similarly the mediation layer), it holds because $[\square_i^e]^e \equiv \text{TypeVal type}_i \text{TypeErr}_i$ has type type_j which is convertible to $\llbracket \square_j^e \rrbracket^e$. For the pure layer, it holds trivially because $[\square_i^p]^p \equiv \square_i$. For all the new constants in RETT that have not been considered in [Pédrot and Tabareau 2018], such as modalities and the parametricity predicate, one only has to check that their translations type check. The last property is a direct application of typing soundness. $\square$

The parametric translation for terms and types that live in the mediation layer is also sound.

THEOREM 4.3 (PARAMETRICITY SOUNDNESS). *The two following properties hold.*

- *If $M \equiv N$ then $[M]_{\varepsilon} \equiv [N]_{\varepsilon}$.*
- *If $\Gamma \vdash M : A : \square_i^m$ then $\llbracket \Gamma \rrbracket_{\varepsilon} \vdash [M]_{\varepsilon} : \llbracket A \rrbracket_{\varepsilon} [M]^m$.*

PROOF. By induction on the derivation. The new typing rules in RETT that have not been considered in [Pédrot and Tabareau 2018] are in the definitions that crosscut the different layers. $\square$

The fact that Theorems 4.2 and 4.3 hold on the whole translation of RETT into CIC allows us to automatically lift many metatheoretical properties of CIC to RETT.

The first obvious one is the consistency of the mediation and pure layers.

THEOREM 4.4 (CONSISTENCY). *The pure layer and the mediation layer of RETT are logically consistent.*

PROOF. Theorem 4.2 on the pure layer guarantees that if M inhabits $\perp_p$ in the pure layer, then $[M]^p$ inhabits $[\perp_p]^p \equiv \perp$ in CIC. In the mediation layer, if M inhabits $\perp_m$, then by Theorem 4.3, $[M]_\varepsilon$ inhabits $[\perp_m]_\varepsilon$. $[M]^m \equiv \perp_\varepsilon [M]^m$, which is equivalent to $\perp$ because $\perp_\varepsilon$ has no constructor. $\square$

The pure and mediation layers of RETT also enjoy a form of canonicity. Canonicity (for booleans) in CIC says that any closed term of type $\mathbb{B}$ is convertible to either true or false. In RETT, we do not know if canonicity holds for the standard conversion (that is, the equational theory arising from the usual rules of CIC together with the additional rules provided for RETT combinators), because this result amounts to showing the completeness of computational laws with respect to the new constants introduced. However, we can prove canonicity for a stronger form of conversion, namely the conversion induced by the translation in CIC, which is complete by definition:

$$M \equiv_{\square} N := [M] \equiv [N].$$

THEOREM 4.5 (CANONICITY). *The pure layer and the mediation layer of RETT enjoy canonicity for $\equiv_{\square}$.*

PROOF. By Theorem 4.2, any closed term M of type $\mathbb{B}$ in the pure layer gives rise to a closed term $[M]^p$ of type $\mathbb{B}$ in CIC. By canonicity, $[M]^p$ is either convertible to true or false, but then as $[\text{true}]^p \equiv \text{true}$ (and similarly for false), we have that $M \equiv_{\square} \text{true}$ or $M \equiv_{\square} \text{false}$ in RETT.

For the mediation layer, the situation is slightly more complicated. Theorem 4.3 guarantees that any closed term M of type $\mathbb{B}$ in the mediation layer gives rise to a closed term $[M]_\varepsilon$ of type $\mathbb{B}_\varepsilon$ in CIC. By canonicity of $\mathbb{B}_\varepsilon$ in CIC, this means that $[M]_\varepsilon$ is convertible to either true_ε or false_ε , and so $[M]^m$ is convertible to either $\text{true}^\bullet$ or $\text{false}^\bullet$. The property follows from the fact that $[\text{true}]^m \equiv \text{true}^\bullet$ (and similarly for false). $\square$

Beyond CIC. While we have formulated RETT as an extension of CIC, it is also interesting to consider extensions of CIC with certain axioms, such as function extensionality.

Interestingly, the translation of RETT satisfies a conservativity result for the pure layer, which states that every axiom that is compatible with CIC is also compatible with the pure layer. To state this theorem, we need to consider the trivial embedding $[-]_{\text{CIC}}$ of CIC into RETT, which is defined by congruence everywhere but for the universes, with $[\square_i]_{\text{CIC}} := \square_i^p$.

THEOREM 4.6 (CONSERVATIVITY OF THE PURE LAYER). *Given an axiom Ax , if $\text{CIC} + Ax$ is consistent, then the pure layer of RETT + $[Ax]_{\text{CIC}}$ is also consistent.*

Finally, we observe that this conservativity result does *not* hold for the mediation layer. For instance, function extensionality can be negated in the mediation layer. This has already been observed in [Pédrot and Tabareau \[2018\]](#).

THEOREM 4.7 (NEGATION OF FUNCTION EXTENSIONALITY [PÉDROT AND TABAREAU \[2018\]](#)). *Function extensionality is not valid in $\square^m$.*

PROOF. The two functions $\lambda x : \top. x$ and $\lambda x : \top. \text{tt}$ can be distinguished in the mediation layer, because we can construct a parametric predicate that observes that the former re-raises exceptions, while the latter is constant and does not. $\square$

In particular, the previous theorem shows that the mediation layer does not preserve univalence [[Univalent Foundations Project 2013](#)], which (coarsely) states that two equivalent types are equal. However, it can be shown using the translation that the mediation layer preserves Uniqueness of Identity Proof (UIP). This means that the mediation layer has to be considered with care when seen as a logical layer.

5 IMPLEMENTATION IN COQ

The translation of RETT into CIC can be seen as a compilation phase that extends the theory of Coq using a plugin, similarly to other syntactic translations [Jaber et al. 2016; Pédro and Tabareau 2017, 2018]. By construction, this does not require any modification to the Coq kernel. Only two additional properties need to be trusted when using the plugin:

- First, that the soundness theorems 4.2 and 4.3 hold.
- Second, that the plugin implements the two translations correctly.

Due to the intrinsic nature of syntactic models, even this additional trust is relative. Soundness failure would merely result in a Coq type error, as the translated terms are still checked by the (unmodified) kernel. Bogus mistranslation is more worrisome, because the compiled term could be unrelated to what the user had in mind. Thankfully, this kind of issue is very similar in spirit to Pollack-inconsistency [Wiedijk 2012], and can be worked around in the same way. Namely, the user can always pierce through the plugin abstraction layer, and scrutinize the translated proof term to check that it indeed corresponds to what was expected.

The specific problem to be solved for RETT is that it has three different hierarchies of universes, which is not the case in Coq. However, we can define a subtheory of RETT that only mentions $\square_i^e$ and $\square_i^p$. By further assuming that $\square^p$ is impredicative, we can implement a plugin that adds exceptions to Coq, where the universe hierarchy **Type** is interpreted as the hierarchy of exceptional types $\square^e$, and the impredicative universe **Prop** is interpreted as the (impredicative) universe of pure types $\square^p$.⁴ The mediation layer $\square^m$ is omitted, but its internal parametricity predicate $\mathcal{P}$ is realized as a Coq type class [Sozeau and Oury 2008].

In this section, we first present how to implement the plugin in Coq. We then explain how we use the type class mechanism to represent parametricity. Finally, we come back to the examples introduced in Section 1.

The code of the example of this section can be found in the file `list_theorem.v` of the anonymous supplementary material.

5.1 CoqRETT: RETT as a Coq plugin

We define a Coq plugin that implements the translation described in Section 4 and provides new constructors in Coq, giving them meaning through the translation. Currently, the plugin does not instrument all the constants of CoqRETT so the user needs to define those constants explicitly.

To define a new constant $C:A$ in CoqRETT, we need to provide a constant of the translation of $[C]:[A]$ in Coq. This is done using the command

```
Effect Definition C : A.
(** definition of [C] **)
Defined.
```

When working inside CoqRETT as a source theory, new definitions can then be introduced as in standard Coq.

The basic new primitives that are available in CoqRETT are the type of exceptions and the function that raises an exception at any type.

```
Definition Exception : Type.
Definition raise : ∀ A : Type, Exception → A.
```

⁴Note that impredicativity is orthogonal to purity; we just exploit the existence of two separate universe hierarchies in Coq.

Note that because of the cumulativity of universes in Coq, there is no way to prevent a user to raise an exception in `Prop`, as `Prop` is a subtype of `Type`. However, translating such an exception will produce a translation error. This means that the correctness of a proof in CoqRETT is not guaranteed only by typechecking the proof, but by additionally typechecking the translation of the proof.⁵

When we define an inductive type in `Type`, for instance lists

```
Inductive list (A : Type) : Type :=
| nil : list A
| cons : A → list A → list A
```

we generate the standard eliminator on `Type`, but we can also provide the catch eliminator on `Type` giving it meaning with the translation:

```
Effect Definition list_catch : ∀ A (P : list A → Type),
P nil → (∀ (a : A) (l : list A), P l → P (a :: l)) → (∀ e, P (raise A e)) → ∀ l : list A, P l.
```

Similarly, we can define the corresponding eliminator `list_catch_prop` in `Prop`.

After going through the translation phase, the computation laws of `list_catch` can be proven by reflexivity, which means that they are indeed definitionally valid in RETT. For instance,

```
Effect Definition list_catch_nil_eq : ∀ A (P : list A → Type) Pnil Pcons Praise,
list_catch A P Pnil Pcons Praise nil = Pnil.
```

Proof.

`reflexivity.`

Defined.

However, these laws can only be proven for *propositional equality* in CoqRETT. This is a limitation in the usability of the plugin, due to the fact that a Coq plugin cannot extend the conversion of the Coq kernel. Indeed, RETT extensions are defined as axioms in the CoqRETT surface language, and computing with them would require a way to make these equalities *definitional*.⁶ Therefore, explicit rewriting with these equalities is necessary when staying in the source theory.

Using the catch eliminator, it is already possible to prove internally in CoqRETT, for instance, that the empty list `nil A` can be discriminated from `raise (list A) e`.

```
Definition nil_not_raise : ∀ A e, nil A ≠ raise e.
```

Proof.

```
intros A e.
assert (H:∀ l', nil A = l' → list_catch _ _ True (fun _ _ => False) (fun _ => False) l').
{ intros l' eq. destruct eq. rewrite list_catch_nil_eq. exact I. }
intro eq. specialize (H (raise e) eq). rewrite list_catch_raise_eq in H. exact H.
```

Defined.

As usual in Coq, the proof relies on dependent elimination. It starts by generalizing `nil A ≠ raise e` to $\forall l', \text{nil } A = l' \rightarrow \text{list_catch } _ _ \text{True } (\text{fun } _ _ \Rightarrow \text{False}) (\text{fun } _ \Rightarrow \text{False}) l'$. Of course, when `l'` is `raise e`, this is the same proposition, but generalizing it allows us to do elimination on the proof of equality, which tells us that we must be in the `nil A` case. Note that in the proof,

⁵We could make this more transparent to the user by instrumenting typechecking to perform both standard typechecking of the term and of its translation.

⁶Extending the reach of the plugin architecture of Coq to support new conversion rules is an interesting perspective, although far beyond the scope of this work.

we need to do explicit rewriting with `list_catch_nil_eq` because `list_catch` does not compute. Once the generalization is proven, the property follows by specialization and rewriting.

5.2 $\mathcal{P}$ as a Type Class

The internal parametricity predicate $\mathcal{P}$ of the mediation layer is realized in CoqRETT as a type class. Indeed, because not every type in $\square^e$ is of the form $\{A\}_e^m$ for some A in $\square^m$, in general, the parametricity predicate is not defined on every type.

The Param type class is used to denote parametric terms.

```
Class Param (A : Type) : Type := { param : A → Prop }
```

The plugin automatically generates instances of the Param type class for inductive types using the parametric translation. Note that we do not make any distinction between the type of lists that comes from the mediation layer and the type of lists in the exceptional layer. This is valid because the two types are isomorphic (Lemma 3.4) and even translated to the same type (Lemma 4.1).

In order to be able to exploit the parametricity predicate for reasoning on inductive types, we add a class ParamInd, which provides the `param_correct` property for instances of Param on inductive types.

```
Class ParamInd (A : Type) {Param A} : Type :=
  { param_correct : ∀ e : Exception, param (raise A e) → False }
```

This property captures the reasoning principle that when a term is parametric, it cannot be an exception. This corresponds to Lemma 3.5 in RETT, here expressed as a primitive notion.

We can recover the eliminator restricted to parametric terms defined in Section 3.5 by using induction on the catch eliminator and the `param_correct` property. In the case of recursive inductive types, we first need to provide an inversion principle that the parametricity of a term implies the parametricity of its subterms, which in the case of lists amounts to

Effect Definition `param_list_cons` : $\forall A \ a \ (l : \text{list } A), \text{param}(\text{cons } a \ l) \rightarrow \text{param } l$.

Then, the eliminator restricted to parametric terms can be defined as

```
Definition list_ind : ∀ A (P : list A → Prop),
  P nil → (∀ (a : A) (l : list A), P l → P (a :: l)) → ∀ l : list A, param l → P l.
```

Proof.

```
intros A P Pnil Pcons l; induction l using list_catch_prop.
- intro. exact Pnil.
- intros param_al. exact (Pcons a l (IHl (param_list_cons _ _ param_al))).
- intros param_e. destruct (param_correct e param_e).
```

Defined.

5.3 Back to Examples

Let us come back to the examples of Section 1, explaining how to state and prove the described results. We insist that all the reasoning that follows is done directly in CoqRETT (as opposed to in Coq over the results of the translation). In the examples we fix the exception type to strings.

Tail of Non-Empty Lists. We can prove in CoqRETT that the exception-raising tail function

```
Definition tail {A} (l : list A) : list A :=
  list_rect (fun _ => list A) (raise "error: empty list") (fun _ l _ => l) l.
```

does not raise an exception when applied to a non-empty list.

To prove the tail property, we first need to establish that, when an integer is provably bigger than another integer, it cannot be an exception. This is because there is no constructor for exceptions in the definition of $\leq$ —comparison is an inductively-defined predicate in `Prop`, and is therefore pure.

Definition `raise_not_leq` : $\forall (n:\mathbb{N})\ e, n \leq \text{raise } e \rightarrow \text{False}$.

This discrimination property (directly induced by the catch eliminator) allows us to prove the non-failing behavior of tail on non-empty lists.

Definition `non_empty_list_distinct_tail_error` : $\forall A\ e\ (l:\text{list } A),$
 $\text{length } l > 0 \rightarrow \text{tail } l \neq \text{raise } e$.

Proof.

```
intros A e l; induction l using list_catch_prop; cbn.
- inversion 1.
- intros Hlen eq. apply le_S_n in Hlen. eapply raise_not_leq. rewrite eq in Hlen.
  rewrite list_rect_raise_eq in Hlen. exact Hlen.
- intros Hlen. unfold length in Hlen. rewrite list_rect_raise_eq in Hlen.
  destruct (raise_not_leq _ _ Hlen).
```

Defined.

The proof is quite direct using induction on the catch eliminator and the `raise_not_leq` discrimination property. The only additional reasoning burden is due to the absence of computation rules for catch elimination, as discussed previously.

Head of Non-Empty Lists. Let us now turn to the exception-raising head function

Definition `head` $\{A\} (l:\text{list } A) : A :=$
`list_rect (fun _  $\Rightarrow A$ ) (raise "error: empty list") (fun a _  $\Rightarrow a$ ) l.`

Recall that proving that applying head on a non-empty list does not produce an exception requires a *deep* notion of parametricity for lists. Deep parametricity is necessary to say that not only the shape of the list is non-exceptional, but also that its contained values are non-exceptional. Of course, this notion of deep parametricity only makes sense for element types for which there exists an instance of the `Param` type class. The predicate `list_param_deep` below defines deep parametricity for lists:

Definition `list_param_deep` $A \{H:\text{Param } A\} : \forall (l:\text{list } A), \text{Prop} :=$
`list_catch A (fun _ : list A  $\Rightarrow$  Prop)`
`True`
`(fun (a : A) (_ : list A) (lind : Prop)  $\Rightarrow$  param a  $\wedge$  lind)`
`(fun _ : Exception  $\Rightarrow$  False).`

It uses the catch eliminator, returning `True` in the case of an empty list and `False` in the case of an exception. The difference with the shallow parametricity predicate is in the recursive case `cons a l`: we require both `a` and `l` to be parametric, the former using the instance of `Param` in hypothesis, and the latter with a recursive call to `list_param_deep`.

With this extra assumption on the list, we can now state and prove the correctness property of head for non-empty lists.

Definition `head_empty_list_no_error` : $\forall A \{H:\text{Param } A\} e (l:\text{list } A),$
 $\text{length } l > 0 \rightarrow \text{list_param_deep } l \rightarrow \text{head } l \neq \text{raise } e$.

Proof.

```

intros A A_param e l. induction l using list_catch_prop.
- inversion 1.
- intros Hlen Hl. unfold list_param_deep in Hl.
  rewrite list_catch_cons_eq in Hl. cbn in *.
  destruct Hl as [Ha _]. intro eq. rewrite eq in Ha. apply (param_correct e Ha).
- intros. unfold length in H. rewrite list_rect_raise_eq in H. compute in H.
  destruct (raise_not_leq _ _ H).
Defined.

```

The proof is again quite direct using induction on the catch eliminator and the `raise_not_leq` discrimination property. The only extra reasoning is in the case the list is actually of the form `cons a l`, where we use the deep parametricity of the list to know that `a` is actually parametric, which by `param_correct` means that it cannot be an exception.

6 RELATED WORK

This work relates to the large body of work on integrating effects and dependent types. Hoare Type Theory (HTT) [Nanevski et al. 2008], used in particular in the Ynot project [Chlipala et al. 2009], is realized as an axiomatic extension of Coq with effects encapsulated in a Hoare monad. HTT does not address the main challenge of effectful terms at the type level because it essentially only supports proving in Coq properties on *simply-typed* imperative programs. Dependent ML [Xi and Pfenning 1999] also side-steps the issue by only allowing types to depend on pure terms, namely arithmetic expressions that denote array lengths. The F^* programming language [Swamy et al. 2016] uses a notion of primitive effects including state, exceptions, divergence and IO. Each effect is described through a monadic predicate transformer semantics. The use of monads makes it possible to isolate a pure core dependent language to reason about effectful programs. However, the standard monadic approach [Moggi 1991] does not scale to dependent types, because one cannot provide a dependently-typed version of the bind operation. Idris [Brady 2013] favors algebraic effects instead of monads as an elegant way to combine effects and dependent types, though with the same restrictions. In contrast, RETT supports reasoning about exceptional programs that can make use of the full power of dependent types.

An alternative, and much lower-level way to address the issue is to represent the effectful fragment of the type theory as a deep embedding of the syntax of this fragment inside the theory. This happens commonly in the implementation of compilers in some flavor of type theory, like e.g. CompCert [Leroy et al. 2016] or Cake ML [Kumar et al. 2014]. While this approach is extremely simple and readily available in weak theories such as LF, it is completely oblivious of the advantages of dependent types. That is, the equational rules of the embedded language have to be computed explicitly in the proof, which in turn also requires proving that the properties of the host language are stable under these rules. As such, handling advanced features like higher-order functions is painful, let alone the preservation of typing of the various programs being considered.

RETT builds upon the translation approach to extend type theory non-axiomatically [Boulier et al. 2017]. Internal translations of type theory have a fairly extensive history. Barthe et al. [1999] describe a CPS translation for CC_ω extended with `call/cc`, which does not handle inductive types. A variant of this translation that supports dependent sums using answer-type polymorphism was developed by Bowman et al. [2018]. Jaber et al. [2016] use forcing to define a generic class of internal translations of type theory that only work on a restricted version of dependent elimination. This limitation also applies to the Baclofen type theory [Pédrot and Tabareau 2017]. RETT is an extension of the Exceptional Type Theory (ETT) [Pédrot and Tabareau 2018], which was the first *complete* internal translation of CIC that adds a specific effect. As discussed earlier, ETT does not

support consistent reasoning about exceptional terms; RETT addresses this limitation through a layered universe architecture with modalities. Both ETT and RETT rely on the internal translation presentation of parametricity of [Bernardy and Lasson \[2011\]](#) in order to impose observational purity on exceptional terms.

A promising venue to reconcile dependent types and effects is to study dependent variants of call-by-push-value (CBPV) [[Levy 2001](#)], as recently done by [Ahman et al. \[2016\]](#) and [Vákár \[2015\]](#). While the CBPV setting can accommodate any effect described in monadic style, these approaches also need to impose a purity restriction for dependency. In contrast, the separation of the pure, mediation, and exceptional layers in RETT makes it possible to isolate restrictions to specific layers, allowing for instance the exceptional layer to freely mix effects and dependencies.

Finally, the Zombie language [[Casinghino et al. 2014](#)] combines proofs and potentially-diverging programs by clearly separating two fragments of the language. This separation is not unlike the layers of RETT, although Zombie does not make it possible to consistently reason about effectful terms. RETT provides solid type-theoretic foundations that can inform the design of similar practical dependently-typed programming languages, as illustrated by the design of the CoqRETT instantiation. In particular, the need for a mediation layer from which the parametricity predicate can be obtained is a key novelty of this work.

7 CONCLUSION AND FUTURE WORK

The Reasonably Exceptional Type Theory (RETT) supports consistent reasoning about exceptional programs in a full dependently-typed setting. As such, it promises to alleviate the task of developing and proving properties about programs that are inherently partial, as well as easing the interoperability between pure type theories used in proof assistants, and mainstream impure functional languages like OCaml and Haskell.

A key element of RETT is its integration of three universe hierarchies, clearly separating the pure and exceptional types, and introducing a mediation layer in order to allow both to interact in a sound manner. We believe this general approach could be beneficial in order to integrate other effects into type theories, sacrificing neither consistency nor modularity.

If this turns out to work for more general effects, it would also mean that it would be possible to extend type theory with *à la carte* effect systems. More precisely, for every single effect being considered, there would be a corresponding universe hierarchy, together with elimination principles that would provide ways to communicate between those different worlds. Such a presentation would be compatible with the current implementation of proof assistants such as Coq, and it would be easy to cherry-pick the particular subsystem one would like to work in.

Finally, the instantiation and implementation of RETT in Coq reveals the interest of a more powerful extension mechanism that would allow some selected propositional equalities to be treated definitionally. Currently, the plugin forces one to rely on explicit rewriting when using hand-defined RETT primitives, which is a major practical hurdle. Recent work on so-called *rewrite rules* [[Cockx and Abel 2016](#)] suggests that the full RETT theory can be emulated, definitional equations included, with a relatively self-contained extension of the Coq kernel. With such an extension, the plugin would be turned into a tiny shell generating the axioms induced by the translation with their associated rewrite rules. An alternative, but much more invasive solution would be to implement RETT directly in the Coq kernel. While not outright impossible, this would represent a massive amount of work, conflicting with other kinds of extensions like univalence.

REFERENCES

- Danel Ahman, Neil Ghani, and Gordon D. Plotkin. 2016. Dependent Types and Fibred Computational Effects. In *19th International Conference on Foundations of Software Science and Computation Structures*. Springer Berlin Heidelberg, Eindhoven, The Netherlands, 36–54. DOI : http://dx.doi.org/10.1007/978-3-662-49630-5_3
- Gilles Barthe, John Hatcliff, and Morten Heine B. Sørensen. 1999. CPS Translations and Applications: The Cube and Beyond. *Higher Order Symbol. Comput.* 12, 2 (Sept. 1999), 125–170. DOI : <http://dx.doi.org/10.1023/A:1010000206149>
- Jean-Philippe Bernardy and Marc Lasson. 2011. Realizability and Parametricity in Pure Type Systems. In *Foundations of Software Science and Computational Structures*, Vol. 6604. Saarbrücken, Germany, 108–122. DOI : <http://dx.doi.org/10.1007/978-3-642-19805-2>
- Yves Bertot and Pierre Castéran. 2004. *Interactive Theorem Proving and Program Development*.
- Simon Boulrier, Pierre-Marie Pédrot, and Nicolas Tabareau. 2017. The next 700 syntactical models of type theory. In *Proceedings of the 6th ACM SIGPLAN Conference on Certified Programs and Proofs, CPP 2017, Paris, France, January 16-17, 2017*. 182–194. DOI : <http://dx.doi.org/10.1145/3018610.3018620>
- William Bowman, Youyou Cong, Nick Rioux, and Amal Ahmed. 2018. Type-Preserving CPS Translation of Σ and Π Types is Not Not Possible. In *Proceedings of the 45th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '18)*. ACM, New York, NY, USA.
- Edwin Brady. 2013. Idris, a general-purpose dependently typed programming language: Design and implementation. *Journal of Functional Programming* 23, 05 (2013), 552–593.
- Joachim Breitner, Antal Spector-Zabusky, Yao Li, Christine Rizkallah, John Wiegley, and Stephanie Weirich. 2018. Ready, Set, Verify! Applying hs-to-coq to Real-World Haskell Code (Experience Report). *Proceedings of the ACM on Programming Languages* 2, ICFP (Sept. 2018), 89:1–89:16.
- Chris Casinghino, Vilhelm Sjöberg, and Stephanie Weirich. 2014. Combining Proofs and Programs in a Dependently Typed Language. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '14)*. ACM, New York, NY, USA, 33–45. DOI : <http://dx.doi.org/10.1145/2535838.2535883>
- Adam Chlipala, Gregory Malecha, Greg Morrisett, Avraham Shinnar, and Ryan Wisnesky. 2009. Effective Interactive Proofs for Higher-order Imperative Programs. In *Proceedings of the 14th ACM SIGPLAN International Conference on Functional Programming (ICFP '09)*. ACM, New York, NY, USA, 79–90. DOI : <http://dx.doi.org/10.1145/1596550.1596565>
- Jesper Cockx and Andreas Abel. 2016. Sprinkles of Extensionality for Your Vanilla Type Theory. In *22nd International Conference on Types for Proofs and Programs (TYPES 2016)*.
- Thierry Coquand and Gérard P. Huet. 1988. The Calculus of Constructions. *Inf. Comput.* 76, 2/3 (1988), 95–120. DOI : [http://dx.doi.org/10.1016/0890-5401\(88\)90005-3](http://dx.doi.org/10.1016/0890-5401(88)90005-3)
- Guilhem Jaber, Gabriel Lewertowski, Pierre-Marie Pédrot, Matthieu Sozeau, and Nicolas Tabareau. 2016. The Definitional Side of the Forcing. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016*. 367–376. DOI : <http://dx.doi.org/10.1145/2933575.2935320>
- Ramana Kumar, Magnus O. Myreen, Michael Norrish, and Scott Owens. 2014. CakeML: A Verified Implementation of ML. In *Principles of Programming Languages (POPL)*. ACM Press, 179–191. DOI : <http://dx.doi.org/10.1145/2535838.2535841>
- Xavier Leroy, Sandrine Blazy, Daniel Kästner, Bernhard Schommer, Markus Pister, and Christian Ferdinand. 2016. CompCert – A Formally Verified Optimizing Compiler. In *ERTS 2016: Embedded Real Time Software and Systems*. SEE.
- Pierre Letouzey. 2004. *Programmation fonctionnelle certifiée : l'extraction de programmes dans l'assistant Coq*. Ph.D. Dissertation. Université Paris XI.
- Paul Blain Levy. 2001. *Call-by-push-value*. Ph.D. Dissertation. Queen Mary, University of London.
- Assia Mahboubi and Enrico Tassi. 2008. *Mathematical Components*.
- Eugenio Moggi. 1991. Notions of Computation and Monads. *Information and Computation* 93, 1 (July 1991), 55–92. DOI : [http://dx.doi.org/10.1016/0890-5401\(91\)90052-4](http://dx.doi.org/10.1016/0890-5401(91)90052-4)
- Aleksandar Nanevski, Greg Morrisett, and Lars Birkedal. 2008. Hoare type theory, polymorphism and separation. *Journal of Functional Programming* 18, 5-6 (2008), 865–911. DOI : <http://dx.doi.org/10.1017/S0956796808006953>
- Ulf Norell. 2009. Dependently Typed Programming in Agda. In *Advanced Functional Programming (AFP 2008) (LNCS)*, Vol. 5832. Springer, 230–266.
- Pierre-Marie Pédrot and Nicolas Tabareau. 2017. An effectful way to eliminate addiction to dependence. In *32nd Annual Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017*. 1–12. DOI : <http://dx.doi.org/10.1109/LICS.2017.8005113>
- Pierre-Marie Pédrot and Nicolas Tabareau. 2018. Failure is Not an Option - An Exceptional Type Theory. In *Proceedings of the 27th European Symposium on Programming Languages and Systems (ESOP 2018)*, Amal Ahmed (Ed.), Vol. 10801. Thessaloniki, Greece, 245–271.
- The Coq Development Team. 2019. The Coq Proof Assistant Reference Manual. (2019). <https://coq.inria.fr/refman>
- John C. Reynolds. 1983. Types, Abstraction and Parametric Polymorphism.. In *IFIP Congress (2002-01-03)*. 513–523.

- Matthieu Sozeau and Nicolas Oury. 2008. First-Class Type Classes. In *Proceedings of the 21st International Conference on Theorem Proving in Higher-Order Logics*. Montreal, Canada, 278–293.
- Nikhil Swamy, Cătălin Hrițcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean-Karim Zinzindohoue, and Santiago Zanella-Béguelin. 2016. Dependent Types and Multi-Monadic Effects in F^* . In *43rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. ACM, 256–270. <https://www.fstar-lang.org/papers/mumon/>
- Éric Tanter and Nicolas Tabareau. 2015. Gradual Certified Programming in Coq. In *Proceedings of the 11th ACM Dynamic Languages Symposium (DLS 2015)*. ACM Press, Pittsburgh, PA, USA, 26–40.
- Univalent Foundations Project. 2013. *Homotopy Type Theory: Univalent Foundations for Mathematics*. <http://homotopytypetheory.org/book>.
- Matthijs Vákár. 2015. A Framework for Dependent Types and Effects. (2015). arXiv:[arXiv:1512.08009](https://arxiv.org/abs/1512.08009) draft.
- Freek Wiedijk. 2012. Pollack-inconsistency. *Electronic Notes in Theoretical Computer Science* 285 (2012), 85 – 100. DOI: <http://dx.doi.org/10.1016/j.entcs.2012.06.008> Proceedings of the 9th International Workshop On User Interfaces for Theorem Provers (UITP10).
- Hongwei Xi and Frank Pfenning. 1999. Dependent Types in Practical Programming. In *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '99)*. ACM, New York, NY, USA, 214–227. DOI: <http://dx.doi.org/10.1145/292540.292560>