

Propositional Equality for Gradual Dependently Typed Programming*

JOSEPH EREMONDI, University of British Columbia, Canada

RONALD GARCIA, University of British Columbia, Canada

ÉRIC TANTER, PLEIAD Lab, Computer Science Department (DCC), University of Chile, Chile

Gradual dependent types can help with the incremental adoption of dependently typed code by providing a principled semantics for *imprecise* types and proofs, where some parts have been omitted. Current theories of gradual dependent types, though, lack a central feature of type theory: propositional equality. Lennon-Bertrand et al. show that, when the reflexive proof *refl* is the only closed value of an equality type, a gradual extension of the Calculus of Inductive Constructions (CIC) with propositional equality violates static observational equivalences. Extensionally-equal functions should be indistinguishable at run time, but they can be distinguished using a combination of equality and type imprecision.

This work presents a gradual dependently typed language that supports propositional equality. We avoid the above issues by devising an equality type of which *refl* is not the only closed inhabitant. Instead, each equality proof is accompanied by a term that is at least as precise as the equated terms, acting as a witness of their plausible equality. These witnesses track partial type information as a program runs, raising errors when that information shows that two equated terms are undeniably inconsistent. Composition of type information is internalized as a construct of the language, and is deferred for function bodies whose evaluation is blocked by variables. We thus ensure that extensionally-equal functions compose without error, thereby preventing contexts from distinguishing them. We describe the challenges of designing consistency and precision relations for this system, along with solutions to these challenges. Finally, we prove important metatheory: type safety, conservative embedding of CIC, weak canonicity, and the gradual guarantees of Siek et al., which ensure that reducing a program's precision introduces no new static or dynamic errors.

CCS Concepts: • **Theory of computation** → **Type structures; Program semantics.**

Additional Key Words and Phrases: dependent types, gradual types, propositional equality

ACM Reference Format:

Joseph Eremondi, Ronald Garcia, and Éric Tanter. 2022. Propositional Equality for Gradual Dependently Typed Programming. *Proc. ACM Program. Lang.* 6, ICFP, Article 96 (August 2022), 29 pages. <https://doi.org/10.1145/3547627>

1 INTRODUCTION

Gradual dependent types relax the discipline of dependent types, so that programmers can write, type check and run programs with partial type information and omit yet-to-be-devised terms or proofs. These capabilities have the potential to help migrate code from non-dependently typed languages, and reduce the learning curve for newcomers to this rich but complex type discipline.

*This work is partially funded by CONICYT FONDECYT Regular Project 1190058. This work is partially funded by an NSERC Discovery Grant.

Authors' addresses: Joseph Eremondi, Department of Computer Science, University of British Columbia, Canada, jeremond@cs.ubc.ca; Ronald Garcia, Department of Computer Science, University of British Columbia, Canada, rxg@cs.ubc.ca; Éric Tanter, PLEIAD Lab, Computer Science Department (DCC), University of Chile, Santiago, Chile, etanter@dcc.uchile.cl.



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

© 2022 Copyright held by the owner/author(s).

2475-1421/2022/8-ART96

<https://doi.org/10.1145/3547627>

Gradual languages [Siek and Taha 2006] check programs against the type information statically available, comparing types via *consistency* \cong , i.e., equality up to missing type information. Static checks skipped due to partial type information are instead performed at run time when the actual values are known. Programs missing type information might not fail, nor are errors deferred until an unsafe operation is attempted. Rather, partial type information is exploited at run time, so an error occurs when a computation’s result has a type incompatible with the context in which it arises. Gradual dependent types let programmers use type-driven programming with holes [Brady 2017] while running code and executing tests, even when missing parts of types, terms or proofs.

However, existing gradual dependent languages do not support *propositional equality* [Martin-Löf 1982]. The propositional equality type $t_1 =_T t_2$ expresses that t_1 and t_2 are equal inhabitants of type T . Its only constructor is `refl`: $t =_T t$, the proof that every term is equal to itself. Equality is useful for practical dependently typed programming, since it lets a function express pre- and post-conditions by taking or returning equality proofs. Likewise, a programmer can use an equality proof to rewrite the type of the expression they are trying to produce. Propositional equality even lays a path to support GADT-style inductive families, since constructors with different return types can be encoded with non-indexed inductive types and propositional equality [McBride 2000].

Limited means of representing and reasoning about equality have been used in existing gradual languages. GCIC [Lennon-Bertrand et al. 2022] supports *decidable equality* (see §2.2), where a type is computed by pattern-matching on the equated terms. Gradual Refinement Types [Lehmann and Tanter 2017] support first-order constraints in linear integer arithmetic. In contrast, propositional equality is general and lightweight: it works for every type, provides its own construction and elimination principles, and can be used with quantifiers or higher order functions.

Until now, the challenge with gradual equality has been propagating and enforcing equality constraints at run time. The problem is that equated terms may contain functions or dependent function types, both of which bind variables. For example, $(\lambda x. x + 0) =_{\mathbb{N}} (\lambda x. x)$ and $((x : \mathbb{N}) \rightarrow \text{Vec } \mathbb{N} (x + 0)) =_{\text{Type}} ((x : \mathbb{N}) \rightarrow \text{Vec } \mathbb{N} x)$ are well-formed types. Extensionally equality of functions (up to partial information) is undecidable. Comparing functions syntactically, by directly comparing bound variables, is decidable. Such a notion works for compile-time consistency checks, but would be problematic during run-time checks, since it destroys static reasoning principles. Observationally equivalent terms in the static language would be distinguishable in the gradual language, e.g., replacing $\lambda x. x + x$ with $\lambda x. 2 * x$ can cause new dynamic errors. Lennon-Bertrand et al. [2022] show that when `refl` is the only static constructor for propositional equality, its inclusion in a gradual language violates static equivalences. Moreover, code that compares bound variables cannot be easily compiled, since every function now needs a syntactic representation.

This paper presents the language GEq (pronounced “geek”), which adds propositional equality to GCIC, allowing `=`, `refl` and `J` to be used like in the Calculus of Inductive Constructions (CIC), but with a dynamic semantics that is meaningful for gradual types. Our **key insight** is to **represent an equality proof using a witness that captures equality constraints discovered at run time**. Taking inspiration from evidence in Abstracting Gradual Typing (AGT) [Garcia et al. 2016] and middle-types in threesomes [Siek and Wadler 2010], we represent witnesses with a term that is as precise as both equated terms. As a consequence, `refl` and `?` are *not* the only inhabitants of the equality type, avoiding the above impossibility result. Our **contributions** are as follows:

- We demonstrate how equality proofs between imprecise terms are useful for discovering bugs in programs and for guiding the development of static proofs (§2);
- We extend GCIC with propositional equality (§4) by typing equality using consistency witnesses between terms (§4.2). We give operational semantics via a cast calculus, where the eliminator for equality uses casts going through the result type given by the witness (§4.3). To combine witnesses

when casting between equality types, we add witness composition directly as a construct in GEq (§4.4). This operator delays the comparison of neutral terms until their variables are bound to values, so composing statically-equivalent functions does not raise an error;

- We prove type safety, conservative extension of CIC, weak canonicity, and the gradual guarantees for GEq (§5), so imprecision never causes stuck states or new (static or dynamic) errors, and GEq rejects ill-typed CIC programs. Like Siek and Chen [2021], our proofs are parameterized over definitions of consistency and precision, revealing sufficient properties to prove the theorems;
- We define precision and consistency for the cast calculus (§6), showing that they fulfill the previously-identified properties. We separate *static consistency*, whether a term of some type can be used in a given context, from *dynamic consistency*, whether two terms compose without error. These coincide for non-dependent gradual languages, but in GEq they must be separated to respect static equivalences while still rejecting ill-typed static programs.

Section 3 reviews GCIC and its cast calculus CastCIC, upon which the other sections build. Section 7 discusses extensions enabled by GEq's features, along with related and future work.

2 SETTING THE STAGE

2.1 Programming vs. Proving and the Gradual Guarantees

Though programming and proving are connected by the Curry-Howard correspondence, the language features best supporting each task differ. Our focus is dependently typed programming: we consider GEq as a model of a programming language rather than as a type theory for mechanizing mathematics. Nevertheless, we prove important metatheory about GEq that may aid in the development of future gradual type theories.

One goal with GEq is proving the *gradual guarantees* of Siek et al. [2015b], which state that a reduction in precision introduces no new static or dynamic errors. These guarantees are useful for programming because of the contrapositive: if a program has a type error, adding more type information does not remove the error. The types are fundamentally inconsistent and must be changed. By contrast, in current implementations of holes either block reduction, causing errors, or block type checking, hiding errors that would otherwise be statically detectable.

2.2 Relationship to Existing Languages

GEq builds off the work of two existing gradual dependent languages, GDTL [Eremondi et al. 2019] and GCIC [Lennon-Bertrand et al. 2022]. Section 7.3 gives a broader discussion of related work.

GDTL is a Gradual Dependently Typed Language with dependent functions, a universe hierarchy, and decidable type checking. It introduced the *imprecise term* ? , which extended gradual typing to allow imprecision not only in types, but type indices and proof terms. GEq inherits ? from GDTL. Since it is based on AGT [Bañados Schwerter et al. 2021; Garcia et al. 2016], GDTL features ideas similar to GEq's witnesses. However, the authors only discuss equality and inductive types as an extension, omitting it from their metatheory. Also, GDTL uses naive syntactic composition, and suffers from the issues we discuss in the introduction: fully static terms that are observationally-equivalent in the static language may have different run-time behavior in the gradual language.

GCIC is a Gradual version of the Calculus of Inductive Constructions (CIC). It uses a cast calculus approach, extending a restricted version of CIC with inductive types but no indexed inductive families or propositional equality. GEq is a direct extension of GCIC. The GCIC authors prove that no gradual language can simultaneously conservatively extend CIC, have strong normalization, and have *graduality*, a strengthening of the gradual guarantees where decreasing then increasing precision produces an equivalent term. The authors give three variants of GCIC, called $\text{GCIC}^{\mathcal{G}}$, $\text{GCIC}^{\mathcal{N}}$ and GCIC^{\uparrow} , which respectively sacrifice one of strong normalization, graduality, and

conservative extension of CIC, while keeping the other two properties. We build off $\text{GCIC}^{\mathcal{G}}$, because, of the three options, sacrificing strong normalization is most palatable for programming. $\text{GCIC}^{\mathcal{N}}$ violates the gradual guarantees, and GCIC^{\uparrow} is too restrictive for practical programming, so we avoid them both. Logical inconsistency and non-terminating proofs are not as detrimental in programming as in mechanized mathematics: type safety is still guaranteed, and errors due to non-terminating proofs are likely to be discovered when a program is run.

GCIC has no dedicated equality type, but decidable equality types are supported. That is, programmers can write a function that takes two elements of a type, and produces a type that is inhabited if and only if they are equal. The programmer must construct, either manually or with tactics, an equality function for each type whose terms they wish to equate, along with the corresponding elimination principle. Most function types have undecidable equality, and hence are unsupported by this method. Also, common functions on equalities cannot be expressed in their most general form with this method, such as such as $\text{cong} : (f : T \rightarrow S) \rightarrow a =_{\top} b \rightarrow f a =_S f b$. Full propositional equality is more convenient for the programmer.

2.3 A Motivating Example: Eagerly Enforcing Specifications

In this section, we motivate our development with examples of how gradual dependent types can catch errors related to the lengths of lists. A guiding principle of our work is that the types the programmer writes should, as much as possible, be treated as specifications to be checked, either statically or dynamically, regardless of whether their enforcement is required for safety.

Throughout the paper, we write static terms using **red sans-serif font**. Terms from the gradual surface language use *green, italic serif font*. The theory is developed using a gradual cast calculus, which we write using **blue, bold serif font**.

A Buggy Quicksort: We begin by showing how gradual types help the migration of a sorting function from a non-dependently typed language to one with dependent types, and how this migration can help identify bugs. Consider a flawed quicksort implementation:

```

sort : List Float  $\rightarrow$  List Float           | List A = Nil | Cons A (List A)
sort Nil = Nil                             | (++) : List A  $\rightarrow$  List A  $\rightarrow$  List A (concatenation)
sort (Cons h t) = (sort (filter (< h) t)) ++ [h] ++ (sort (filter (> h) t))

```

Since $<$ is used instead of \leq , duplicates are erroneously removed from the list. The programmer may have a suspicion that they have made a mistake in their code, or may have observed incorrect behavior while testing. Their dependent type enthusiast friends have repeatedly assured them that dependently typed languages can help eliminate bugs, so they try migrating their code to a dependently typed language with propositional equality.

An approach to reasoning about the correctness of **sort** is to use *Fixed-length Lists*, which we call **FLists**. Dependent pairs and propositional equality allow for a type of lists indexed by their length.¹ Figure 1 shows this type, and how it can be used to express that **sort** should preserve the length of the produced list. Here, **refl** is the **reflexive** proof that **t** is equal to itself. That **sort** function behaves like the non-dependent version, except it must extract the **Lists** from the **FLists** produced by the recursive calls, and produce an **FList** with proof that the length is the same as the input.

At this stage, the programmer must fill the hole `???` by constructing a proof of type $\text{length} (\text{Cons } h \ t) =_{\mathbb{N}} \text{length} (\text{sortLt } ++ [h] ++ \text{sortGt})$. This task is difficult for a newcomer, since they must use associativity of addition and how $\text{length} (\text{list}_1 ++ \text{list}_2) =_{\mathbb{N}} \text{length } \text{list}_1 + \text{length } \text{list}_2$.

¹A more conventional approach would be to use an indexed type family, which we discuss in §7.1.

Moreover, they must prove that $\text{length } t =_{\mathbb{N}} (\text{length } lt) + (\text{length } gt)$, but such a proof is impossible, due to the bug. Even if they had such a proof, they would need to then use the proofs from the recursive calls, plt and pgt , to relate the lengths of lt and gt to the lengths of sortLt and sortGt .

The programmer is now cursing their type-theorist friend. For a non-buggy quick-sort, one could construct the necessary proof, but doing so is difficult, particularly for a newcomer. The type checker does not detect the bug, so it does not inform the programmer that hole cannot be filled, and it cannot say which aspect of the proof is impossible. Also, development has now stopped: the programmer cannot run or test their code without the missing proof.

Gradual Types to the Rescue: GEq lets the programmer run and test `sort` before writing the missing proof, checking (within the limits of decidability) whether any static values could possibly replace imprecise types and proofs. In Fig. 1, replacing the hole with `?`, the *imprecise term*,² yields a complete, well-typed GEq function that can be called, tested, or used in other modules.

The utility of gradual typing is shown in the run-time checks that let us identify bugs in code. The run-time semantics of GEq are defined via type-directed elaboration to a *cast calculus* CastEq, in which all implicit conversions are replaced by explicit casts. During type checking, `?` is elaborated into the CastEq's $?_{n=\mathbb{N}, \text{length } (\text{sortLt} \# [h] \# \text{sortGt})}$ (the least precise term of type $n =_{\mathbb{N}} \text{length } (\text{sortLt} \# [h] \# \text{sortGt})$), which is not a value in CastEq. Instead, it reduces to the consistency witness $n \&_{\mathbb{N}} \text{length } (\text{sortLt} \# [h] \# \text{sortGt})$. The operator $\&$ is the gradual composition operator, which combines information statically known about its operands. Because types depend on terms, composition is not limited to types, but can combine terms of any type. The $\&$ operator is a syntactic construct of CastEq, not a meta-operation like it is in existing literature [Siek et al. 2015a; Siek and Wadler 2010]. Reifying $\&$ into the object language is critical for composing functions (§2.4). Since n is a variable, this composition expression does not reduce further.

We can identify the bug once `sort` is applied to a concrete list. Consider the input `[2.2, 1.1, 3.3, 2.2]`, which elaborates to `[2.2, 1.1, 3.3, 2.2]` in CastEq. Applying `sort` binds $lt := [1.1]$ and $gt := [3.3]$, giving a result list of `[1.1, 2.2, 3.3]`. Then n is 4 and $\text{length } (\text{sortLt} \# [h] \# \text{sortGt})$ is 3, so the witness for the result is the composition $4 \&_{\mathbb{N}} 3$, which reduces to a run-time error.

In a language without dependent types, this bug could be caught with testing or assertions. In GEq, however, dependent types provide a unified means of specifying properties to be checked statically or dynamically. During development, types serve as assertions to be checked dynamically (or statically, if enough information is present). When a program is completed and all uses of `?` have been removed, those same types establish properties that have been statically verified.

Witness Composition: The key to finding the error above was tracking information with witnesses, and combining those witnesses using the composition operator. While that composition was a simple equality check, in general the composed values may be imprecise, and the result is some value that is as precise as both inputs. The information from the witness is used when

```

FList : Type → ℕ → Type
FList A n = ((x : List A) × length x =ℕ n)
-----
sort : (n : ℕ) → FList Float n → FList Float n
sort 0 (Nil, p) = (Nil, p)
sort (1 + n) (Cons h t, p) =
  let lt = (filter (< h) t)
      (sortLt, plt) = sort (length lt) (lt, refl)
      gt = (filter (> h) t)
      (sortGt, pgt) = sort (length gt) (gt, refl)
  in (sortLt # [h] # sortGt, ???)

```

Fig. 1. Sorting Fixed-Length Lists

²Each `?` is actually `?@0`, i.e., annotated with its type's universe level. Our exposition omits levels; we explain them in §3.3.

eliminating an equality proof: when using a witness t_w of $t_1 =_T t_2$ to rewrite a term of type $P(t_1)$ into $P(t_2)$, we first cast to $P(t_w)$, then to $P(t_2)$. For a program with imprecise types or values, the witness retains the information gained by running the program, preventing unsafe execution, and informing the programmer when a counter-example to an imprecise equality is found.

Here we present an example of a bug that is found, not because of safety, but because a remembered constraint was violated. Consider the following functions:

$$\begin{aligned} \text{zip} &: (n : \mathbb{N}) \rightarrow \text{FList } A \ n \rightarrow \text{FList } A \ n \rightarrow \text{FList } (A \times A) \ n \\ \text{take} &: (n : \mathbb{N}) \rightarrow (m : \mathbb{N}) \rightarrow \text{FList } B \ (n + m) \rightarrow \text{FList } B \ n \end{aligned}$$

Here, *zip* takes two lists of exactly the same length, and produces a list of pairs of their elements, while *take* takes a list with at least n elements, and returns the first n elements of that list. Each function constrains the size of its input, so by tracking equality witnesses, we can also track these constraints and detect where they are incompatible. Now consider lists with imprecise types:

$$\text{list}_1 := ([1.1, 2.2], \text{refl}_?) : \text{FList } \text{Float } ? \qquad \text{list}_2 := (\text{Cons } 1.1 \ ?, \text{refl}_?) : \text{FList } \text{Float } ?$$

For list_1 , we are converting a list of length 2 to a fixed-length list of unknown length, since 2 is consistent with $?$. For list_2 , however, the length is truly imprecise, since its tail is the unknown term. We can zip these lists together as $\text{zip } ? \ \text{list}_1 \ \text{list}_2 : \text{FList } \text{Float } ?$, producing another list of unknown length, since recursively applying *zip* to the unknown tail $?$ produces an unknown result. Applying *take* 3 to the result of *zip* is well typed, since the length $?$ is consistent with $3 + ?$, i.e., $\text{take } 3 \ ? \ (\text{zip } ? \ \text{list}_1 \ \text{list}_2) : \text{FList } \text{Float } 3$. However, computing the witness flags an error.

This error represents something deeper than a simple safety check: it detects fundamental inconsistencies in statically-determined propositional equalities. In the absence of equality proofs, the call could run safely: $\text{Cons } (1.1, 3.3) \ ?$ would be a sensible result, having length consistent with 3. The witness composition is not just checking if a list is empty before taking the head, or counting the elements in the list before running *take*. Rather, the information added by *zip*, that the list should have length 2, has been propagated using the list_1 witness and composed with the conflicting information. GEq uses witnesses to enforce imprecise equality constraints at run time.

To understand how GEq detects this mismatch, we look at the result of elaborating to CastEq. Initially, list_2 has $1 + ?_{\mathbb{N}}$ as the witness that $?_{\mathbb{N}}$ is equal to $1 + ?_{\mathbb{N}}$. The result of *zip* has 2 as the witness of equality between $?_{\mathbb{N}}$ and $?_{\mathbb{N}}$, since that is the length of list_1 . This new witness was determined by composition: since $1 + ?_{\mathbb{N}}$ is consistent with 2, this composition succeeds. (Using a Peano representation of naturals, $S(?)$ is consistent with $S(S(0))$) Then, even though *zip*'s result has a type that is consistent with what *take* expects, the run-time type information remembers that *zip* constrained the list to have length 2. The result of *zip* is cast to $\text{FList } \text{Float } (3 + ?_{\mathbb{N}})$, the type expected by *take* 3. The *zip* result has an equality proof of type $1 + ?_{\mathbb{N}} =_{\mathbb{N}} ?_{\mathbb{N}}$, which is cast to type $1 + ?_{\mathbb{N}} =_{\mathbb{N}} 3 + ?_{\mathbb{N}}$. During this cast, the target value $3 + ?_{\mathbb{N}}$ is composed with the witness 2. Despite the imprecision, these values are not consistent, and composition produces an error: no value can replace $?$ to make $S(S(S(?)))$ equal to $S(S(0))$. We detail the semantics enabling this in §4.

With equality witnesses, we achieve more than type safety. From the gradual guarantees, we know the above code cannot possibly be made static by replacing the $?$ uses with static terms. When a witness reduces to an error, the program is equating two terms that are fundamentally not-equal. So the gradual guarantees now inform about equality constraints, in addition to type constraints. These constraints are expressed through types, rather than some external language of assertions.

2.4 Lazily Enforcing Specifications: Function Equalities and Extensionality

Propositional equality is not restricted to first-order values like numbers or to types with decidable equality. In particular, we can form equalities between functions, for which equality is not in general

decidable. The following summarizes how GEq handles propositional equality for functions without encountering the impossibility result of [Lennon-Bertrand et al. \[2022\]](#). Consider the example they use to show the incompatibility between gradual typing and **refl**-based equality:

$$\begin{aligned} id_{\mathbb{N}} &:= (\lambda x. x) : \mathbb{N} \rightarrow \mathbb{N} & add0 &:= (\lambda x. x + 0) : \mathbb{N} \rightarrow \mathbb{N} \\ test &:= \lambda f. \mathbf{J} (_.\mathbb{B}) id_{\mathbb{N}} \text{true } f (refl_{id_{\mathbb{N}}} :: ? :: (id_{\mathbb{N}} =_{\mathbb{N} \rightarrow \mathbb{N}} f)) : (\mathbb{N} \rightarrow \mathbb{N}) \rightarrow \mathbb{B} \end{aligned}$$

Here **J** is the eliminator for equality: we explain it fully in §4.2, but it suffices to know that in this case, it uses a proof of type $id_{\mathbb{N}} =_{\mathbb{N} \rightarrow \mathbb{N}} f$ to rewrite $(\lambda _.\mathbb{B}) id_{\mathbb{N}}$ to $(\lambda _.\mathbb{B}) f$. The form $::$ denotes surface type ascriptions, which are elaborated to casts in the cast calculus. Both types reduce to \mathbb{B} , but **J** only reduces if the equality proof reduces without error.

Since $id_{\mathbb{N}}$ and $add0$ agree on all inputs, they should be observationally equivalent, producing the same result in any context in which we use them. Violating this would mean that the embedding of CIC into GCIC or GEq does not respect function extensionality, i.e., some statically-equivalent terms are distinguishable in the gradual language. [Lennon-Bertrand et al. \[2022\]](#) offer *test* as a context that distinguishes $id_{\mathbb{N}}$ and $add0$. When $id_{\mathbb{N}}$ is given as an argument, casting $refl_{id_{\mathbb{N}}}$ to $?$ then back to $id_{\mathbb{N}} =_{\mathbb{N} \rightarrow \mathbb{N}} id_{\mathbb{N}}$ should produce $refl_{id_{\mathbb{N}}}$. However, when $add0$ is given for f , the cast must fail, since *refl* cannot have type $id_{\mathbb{N}} =_{\mathbb{N} \rightarrow \mathbb{N}} add0$.

The first key to avoiding this inequivalence is the witness-based representation of equality. In CastEq, **refl** is the only constructor for equality, but it takes an argument: the consistency witness for the equated terms. Moreover, it does not require the equated terms to be syntactically identical, only that the witness be at least as precise as both of them. So **refl** can have type $id_{\mathbb{N}} =_{\mathbb{N} \rightarrow \mathbb{N}} add0$.

What witness should be attached to the proof of $id_{\mathbb{N}} =_{\mathbb{N} \rightarrow \mathbb{N}} add0$? The second key feature is the composition operator of CastEq, which builds said equality witness. Elaborating $refl_{id_{\mathbb{N}}}$ creates witness $id_{\mathbb{N}}$. The cast to $id_{\mathbb{N}} =_{\mathbb{N} \rightarrow \mathbb{N}} f$ composes that witness with the destination endpoints, $id_{\mathbb{N}}$ and f , yielding $id_{\mathbb{N}} \&_{\mathbb{N} \rightarrow \mathbb{N}} id_{\mathbb{N}} \&_{\mathbb{N} \rightarrow \mathbb{N}} f$. The semantics of $\&$ reduce this composition to $\lambda x. (x \&_{\mathbb{N}} x \&_{\mathbb{N}} (f x))$, similar to how a higher-order contract applied to a function produces a new function that checks the input and result [[Findler and Felleisen 2002](#)]. Since $\&$ is an operator in the language, the composition does not need to reduce further, but when the function is applied, it continues to reduce. The same holds when we replace f with $id_{\mathbb{N}}$ or $add0$. Section 6 defines precision such that $x \&_{\mathbb{N}} x \&_{\mathbb{N}} f x$ is more precise than both x and $f x$, so the above composition is a valid witness. We define semantics for **J** so that when it is given the equality proof with the above witness, it reduces, so *test* reduces to **true** for both $id_{\mathbb{N}}$ and $add0$.

How can equating these functions be safe, since deferring composition means that we can prove an equality between unequal functions? As we saw with **sort** above, **J** casts through the witness, so when functions are extensionally non-equal, trying to prove equality between their results dynamically fails. Consider instead $subadd1 := (\lambda(x : \mathbb{N}). x - 1 + 1)$. Since $0 - t = 0$ in \mathbb{N} , $subadd1\ 0 = 1$. We can use the witness $\lambda x. (x \&_{\mathbb{N}} x \&_{\mathbb{N}} (subadd1\ x))$ to inhabit $id_{\mathbb{N}} =_{\mathbb{N} \rightarrow \mathbb{N}} subadd1$. However, if we try to use **J** to prove that $id_{\mathbb{N}}\ 0 =_{\mathbb{N}} subadd1\ 0$, the result substitutes 0 for x in the witness, giving $0 \&_{\mathbb{N}} 0 \&_{\mathbb{N}} 1$, which reduces to an error.

The consequence of our approach is that GEq supports a limited form of extensionality. *Neutral terms*, i.e., variables or terms for which reduction is blocked by one or more variables, always compose to a non-error, so we can build a witness capturing the plausibility of equality between them, given partial information. That witness makes an equality proof constructible. Furthermore, any two functions with neutral bodies compose to a non-error. If the functions agree on all inputs, eliminating their equality never fails and the proof of equality can be freely used. If the functions disagree on some input, an error is raised when building a term that relies on the functions producing the same value for said input. Since it is undecidable whether two functions agree on every input, this approach finds a balance between decidability and flexibility.

Static vs. Dynamic Consistency: For non-dependent gradual types, the successful composition of two types usually implies that they are consistent. However, for GEq, two neutrals always compose to a more-precise term. To conservatively extend CIC, all ill-typed (fully-static) CIC programs must be ill-typed in GEq, and an uninhabited CIC type must not be inhabited by any fully-static GEq terms. So GEq cannot have all neutrals consistent, since this would yield a fully-static proof of $(\lambda x \lambda y. x) =_{\mathbb{T} \rightarrow \mathbb{T} \rightarrow \mathbb{T}} (\lambda x \lambda y. y)$.

We resolve this tension with separate static and dynamic notions of consistency (§6). Terms are statically consistent if they are syntactically equal up to α -equivalence, reduction, and occurrences of $?_{\mathbb{T}}$. Terms are dynamically consistent if they compose without error, or equivalently, if there exists a non-error term as precise as both terms. Essentially, terms are dynamically-consistent if they are statically consistent in the non-neutral parts. The type rules for GEq use static consistency. Some pairs of terms are not statically consistent, yet still compose to a non-error term.

To compare static and dynamic consistency, consider the ill-typed CIC term $\text{refl} : x =_{\mathbb{N}} y$. When embedded into GEq, $\text{refl} : x =_{\mathbb{N}} y$ is still ill-typed: the expected type of $x =_{\mathbb{N}} y$ and the actual type of $x =_{\mathbb{N}} x$ are not statically consistent, because the variables x and y are not identical. In CastEq, x and y are neutral, and hence dynamically consistent, meaning $x \&_{\mathbb{N}} y$ witnesses $x =_{\mathbb{N}} y$.

Allowing neutrals to be dynamically consistent does not interfere with conservatively extending CIC. For conservative extension, every ill-typed CIC program should be ill-typed in GEq. In the absence of $?$, pairs of definitionally-unequal CIC terms are statically inconsistent. While GEq gives refl the same type as CIC, CastEq lets refl prove equality for dynamically consistent terms. However, dynamic consistency does not allow CastEq to type ill-typed CIC terms, because CIC programs are elaborated into a subset of CastEq where t only witnesses $t =_{\mathbb{T}} t$ and all casts have the form $\langle \mathbb{T} \leftarrow \mathbb{T} \rangle$. The type $(x : \mathbb{N}) \rightarrow (y : \mathbb{N}) \rightarrow (x =_{\mathbb{N}} y)$ is uninhabited in CIC, and while this type is inhabited in CastEq using witness $x \&_{\mathbb{N}} y$, the use of $\&_{\mathbb{N}}$ puts the witness outside the static fragment of CastEq. The term $x \&_{\mathbb{N}} y$ does not correspond to any typed or ill-typed CIC program.

Static and dynamic consistency let us balance conflicting goals. If all statically inconsistent functions composed to an error, then statically-equivalent terms would not be gradually equivalent, making it harder to reason about program equivalence. Using dynamic consistency during type checking would not conservatively extend CIC. By separating these, we obtain conservative extension, while dynamically respecting all extensional equalities.

3 THE STATIC LANGUAGE AND GCIC

To begin our development, we review the state-of-the-art for handling inductive types in a gradual language. We describe the Bidirectional CIC (BCIC), a modification of CIC whose bidirectional types are convenient for gradual typing [Lennon-Bertrand 2021]. We then describe the gradual surface language GCIC, along with the cast calculus, CastCIC, and a translation from GCIC to CastCIC [Lennon-Bertrand et al. 2022]. Specifically, we use $\text{GCIC}^{\mathcal{G}}$, the variant that satisfies the gradual guarantees and embeds CIC, but sacrifices strong normalization. We discuss options for decidable type checking in §7.2.1. Though GCIC is not a contribution of this paper, we use it as the starting point for our development, making additions to the surface language and cast calculus.

3.1 Bidirectional CIC

3.1.1 Syntax. Figure 2 gives the *bidirectional calculus of inductive constructions* (BCIC) as originally presented by Lennon-Bertrand [2021], though we modify their notation to maximize clarity for GEq's additions. BCIC terms are denoted by metavariables t and T , loosely following the convention that T be reserved for types. Variables are denoted by x, y, z . BCIC has variables, a predicative hierarchy of universes, function types, functions, and applications. Technically, BCIC extends the predicative, non-cumulative fragment of CIC: each function type is in a higher universe than its

$t, T ::= x \mid \mathbf{Type}_i \mid (x : T_1) \rightarrow T_2 \mid \lambda(x : T). t \mid t_1 t_2 \mid C_{@i}(\bar{t}) \mid D^C_{@i}(\bar{t}, \bar{t}') \mid \mathit{ind}_C(t_1, z.T, x.\overline{y.t_2})$	
$h ::= C \mid \mathbf{Type} \mid \Pi$	
$\Gamma \vdash t \Rightarrow T$ (Synthesis) $\Gamma \vdash t \Leftarrow T$ (Checking) $\Gamma \vdash t \Rightarrow_h T$ (Constrained Synthesis)	
$\frac{\Gamma \vdash t \Rightarrow T' \quad T \longrightarrow^* T'' \quad T' \longrightarrow^* T''}{\Gamma \vdash t \Leftarrow T} \quad \frac{(x : T) \in \Gamma}{\Gamma \vdash x \Rightarrow T} \quad \frac{}{\Gamma \vdash \mathbf{Type}_i \Rightarrow \mathbf{Type}_{i+1}}$	
$\frac{\Gamma \vdash T_1 \Rightarrow_{\mathbf{Type}} \mathbf{Type}_i \quad \Gamma, x : T_1 \vdash T_2 \Rightarrow_{\mathbf{Type}} \mathbf{Type}_j}{\Gamma \vdash (x : T_1) \rightarrow T_2 \Rightarrow \mathbf{Type}_{\max(i,j)}} \quad \frac{\Gamma \vdash t_0 \Rightarrow_{\Pi} (x : T_1) \rightarrow T_2 \quad \Gamma \vdash t_1 \Leftarrow T_1}{\Gamma \vdash t_0 t_1 \Rightarrow [t_1/x]T_2} \quad \frac{\Gamma \vdash T_1 \Rightarrow_{\mathbf{Type}} \mathbf{Type}_i \quad \Gamma, x : T_1 \vdash t \Rightarrow T_2}{\Gamma \vdash \lambda(x : T_1). t \Rightarrow (x : T_1) \rightarrow T_2}$	
$\frac{\Gamma \vdash t_k \Leftarrow \mathbf{Params}_k(C, i)[\bar{t}]}{\Gamma \vdash C_{@i}(\bar{t}) \Rightarrow \mathbf{Type}_i} \quad \frac{\Gamma \vdash t_k \Leftarrow \mathbf{Params}_k(C, i)[\bar{t}]}{\Gamma \vdash D^C_{@i}(\bar{t}, \bar{t}') \Rightarrow C_{@i}(\bar{t})} \quad \Gamma \vdash t'_m \Leftarrow \mathbf{Args}_m(C, i, D)[\bar{t}, \bar{t}']$	
$\frac{\Gamma \vdash t_{\text{scrut}} \Rightarrow_C C_{@i}(\overline{t_{\text{par}}}) \quad \Gamma, z : C(\overline{t_{\text{par}}}) \vdash T_P \Rightarrow_{\mathbf{Type}} \mathbf{Type}_j \quad \left(\frac{\Gamma, x_{\text{rec}} : X, \bar{y} : \mathbf{Args}(C, i, D_k)[\overline{t_{\text{par}}}, \bar{y}] \vdash t_{\text{rhs}_k} \Leftarrow [D^C_{@i}(\overline{t_{\text{par}}}, \bar{y})/z]T_P}{\text{where } X := (z : C_{@i}(\overline{t_{\text{par}}})) \rightarrow T_P} \right)^k}{\Gamma \vdash \mathit{ind}_C(t_{\text{scrut}}, z.T_P, x_{\text{rec}}.\overline{y.t_{\text{rhs}}}) \Rightarrow [t_{\text{scrut}}/z]T_P} \quad \frac{\Gamma \vdash t \Rightarrow T \quad T \longrightarrow^* \mathbf{Type}_i}{\Gamma \vdash t \Rightarrow_{\mathbf{Type}} \mathbf{Type}_i}$	
$\frac{\Gamma \vdash t \Rightarrow T \quad T \longrightarrow^* (x : T_1) \rightarrow T_2}{\Gamma \vdash t \Rightarrow_{\Pi} (x : T_1) \rightarrow T_2} \quad \frac{\Gamma \vdash t \Rightarrow T \quad T \longrightarrow^* C_{@i}(\overline{t_p})}{\Gamma \vdash t \Rightarrow_C C_{@i}(\overline{t_p})}$	
$t \rightsquigarrow t'$ (Notions of Reduction) $t \longrightarrow t'$ (Contextual reduction)	
$\frac{(\lambda(x : T). t) t' \rightsquigarrow [t'/x]t \quad t \rightsquigarrow t' \quad C \text{ an arbitrary context}}{C[t] \longrightarrow C[t']}$	
$\mathit{ind}_C(D_k(\overline{t_{\text{par}}}, \overline{t_{\text{arg}}}), z.T_P, x_{\text{rec}}.\overline{y.t_{\text{rhs}}}) \rightsquigarrow [\lambda(x : C(\overline{t_{\text{par}}})) . \mathit{ind}_C(x, z.T, x_{\text{rec}}.\overline{y.t_{\text{rhs}}})/x_{\text{rec}}][\overline{t_{\text{arg}}}/\bar{y}]t_{\text{rhs}_k}$	

Fig. 2. Bidirectional CIC: Syntax, Typing and Semantics

domain and codomain, and there is no subtyping between universe levels. We assume a pre-existing set of inductive type constructors, denoted by the metavariable C , each of which has a fixed set of data constructors D^C . Type and data constructors are annotated with the level of their type, though we omit these annotations when they are not relevant.

To eliminate members of inductive types, a combined form $\mathit{ind}_C(t_1, z.T, x.\overline{y.t_2})$ replaces CIC's `fix` and `match`. This form branches on the scrutinee t_1 and has a parameterized result type T , called the *motive* [McBride 2002], that binds a variable of the scrutinee's type. The branches t_2 correspond to the constructors D^C of C . In each branch, the variables \bar{y} are bound to the arguments to D^C , and x refers to the whole ind_C expression, to facilitate recursion. The ind_C form expresses an induction principle: if each branch produces a result of type T where z is bound to D^C applied to \bar{y} , the elimination has type T where z is bound to the scrutinee t_1 . In essence, ind_C says that if we can build a T for each constructor D^C of $C(\overline{t_{\text{par}}})$, then we can build one for any value of $C(\overline{t_{\text{par}}})$. Normally, a separate check ensures that recursive calls are only made on structurally smaller

$$\boxed{t, T ::= ?_{@i} \mid t :: T}$$

$$\frac{}{\Gamma \vdash ?_{@i} \Rightarrow ?_{@i+1}} \quad \frac{\Gamma \vdash t \Leftarrow T}{\Gamma \vdash t :: T \Rightarrow T} \quad \frac{\Gamma \vdash t \Rightarrow T' \quad T' \cong_{\rightarrow} T}{\Gamma \vdash t \Leftarrow T}$$

$$\frac{\Gamma \vdash t \Rightarrow T \quad T \longrightarrow^* ?_{@i}}{\Gamma \vdash t \Rightarrow_{\Pi} ?_{@i} \rightarrow ?_{@i}} \quad \frac{\Gamma \vdash t \Rightarrow T \quad T \longrightarrow^* ?_{@i}}{\Gamma \vdash t \Rightarrow_C C(\overline{?_{\text{Params}(C,i)}})} \quad \frac{\Gamma \vdash t \Rightarrow T \quad T \longrightarrow^* ?_{@i+1}}{\Gamma \vdash t \Rightarrow_{\text{Type}} \text{Type}_i}$$

Fig. 3. GCIC: Syntax and Typing Lemmas

arguments, but we omit this check, since it is orthogonal to gradual typing and GEq would not be strongly normalizing even with it.

BCIC also uses *head tags*, denoted by h , which act as symbols to specify a type constructor without specifying its arguments. We use these in typing, e.g. for expressing that an applied function must synthesize a function type, even though we do not know what the domain and codomain should be. Tags are also useful in GCIC for defining the least precise type with a given head.

3.1.2 Typing and Semantics. The typing and semantics (Fig. 2) for BCIC resemble the typical presentation of CIC, but typing is divided into synthesis, which produces a type, and checking, which consumes a type. The semantics is given with primitive notions of reduction \rightsquigarrow , contextual stepping \longrightarrow where any sub-term reduces (even under binders), and multi-step contextual reduction \longrightarrow^* , which allows zero or more steps with \longrightarrow . Because function types bind their parameter in the codomain, applications synthesize a type depending on the value of the argument, since it is substituted for x in the codomain type. A term checks against any type that reduces to the same type as its synthesized type, since an application may have produced a type that must be reduced before comparing. Constrained-synthesis, $\Gamma \vdash t \Rightarrow_h T$, generalizes the pattern of synthesizing a type for a term after reducing it to a point that it has the desired head h . Function application has the standard β -reduction rule.

For inductive types, the typing rule establishes that, if T_P is a type parameterized over a value x from the inductive type C , and we can (recursively) build a T_P for each constructor of C , then we can build a T_P for any member of C . Hence $\text{ind}_C()$ form gives an induction principle for C , hence the notation $\text{ind}_C(\dots)$. Inductive types may be parameterized, but each constructor has the same return type. The reduction rule says that an $\text{ind}_C(\dots)$ form given a value $D^C(\dots)$ reduces to the branch corresponding to D^C .

3.2 GCIC: The Surface Language

Figure 3 extends BCIC into GCIC, the Gradual CIC, by adding the *imprecise term* $?_{@i}$, which can be used at any type in universe level i , along with type ascriptions, which were not in BCIC because all forms synthesized types. We use $?_T$ as sugar for $?_{@i}$ when $T : \text{Type}_i$.

Dependent types complicate the typing of GCIC. Because the dynamic semantics of GCIC are defined using a cast calculus, and typing refers to reduction of terms, Lennon-Bertrand et al. [2022] define GCIC typing with cast calculus types. Nevertheless, we can establish lemmas (Fig. 3), phrased like rules, which provide intuition for how GCIC terms are typed against GCIC types, helping GCIC be understood without diving into the details of the cast calculus.

The unknown term $?_{@i}$ synthesizes $?_{@i+1}$ i.e. its type is unknown, one level up in the universe hierarchy. A term checks against any type consistent with its synthesized type, where the relation \cong_{\rightarrow} is understood to mean convertibility up to well-typed occurrences of $?$ (explained fully in §6.1.2). An ascribed term synthesizes the given type if it checks against it, relaxing or tightening the types of gradual terms. Constrained synthesis accounts for terms synthesizing $?_{\text{Type}_i}$, by producing the *germ* (called the *ground type* in non-dependently typed literature). The type germ $_i(h)$ is the

$$\boxed{t, T ::= ?_T \mid \mathcal{U}_T \mid \langle T_2 \Leftarrow T_1 \rangle t}$$

$$\frac{\Gamma \vdash T \Rightarrow_{\text{Type}} \text{Type}_i}{\Gamma \vdash ?_T \Rightarrow T} \quad \frac{\Gamma \vdash T \Rightarrow_{\text{Type}} \text{Type}_i}{\Gamma \vdash \mathcal{U}_T \Rightarrow T} \quad \frac{\Gamma \vdash T_j \Rightarrow_{\text{Type}} \text{Type}_i \text{ for } j \in \{1, 2\} \quad \Gamma \vdash t \Leftarrow T_1}{\Gamma \vdash \langle T_2 \Leftarrow T_1 \rangle t \Rightarrow T_2}$$

Propogation Reductions:

$$\begin{array}{l}
\text{REDPROPFCSTIND(UNK,ERR)} \quad \text{REDPROPFCSTDOWN(UNK,ERR)} \\
\langle C(\bar{t}_2) \Leftarrow C(\bar{t}_1) \rangle ?_{T_1} \rightsquigarrow ?_{C(\bar{t}_2)} \quad \langle T \Leftarrow ?_{\text{Type}_i} \rangle ?_{?_{\text{Type}_i}} \rightsquigarrow ?_T \\
\langle C(\bar{t}_2) \Leftarrow C(\bar{t}_1) \rangle \mathcal{U}_{T_1} \rightsquigarrow \mathcal{U}_{C(\bar{t}_2)} \quad \langle T \Leftarrow ?_{\text{Type}_i} \rangle \mathcal{U}_{?_{\text{Type}_i}} \rightsquigarrow \mathcal{U}_T \\
\\
\text{REDPROPFUN(UNK,ERR)} \quad \text{REDPROPMATCH(UNK,ERR)} \\
?_{(x:T_1) \rightarrow T_2} \rightsquigarrow \lambda(x : T_1). ?_{T_2} \quad \text{ind}_C(?_{C(\bar{t}_{\text{par}})}, z.T_P, x_f.\bar{y}.\bar{t}) \rightsquigarrow ?_{[?_{C(\bar{t}_{\text{par}})}/z]T_P} \\
\mathcal{U}_{(x:T_1) \rightarrow T_2} \rightsquigarrow \lambda(x : T_1). \mathcal{U}_{T_2} \quad \text{ind}_C(\mathcal{U}_{C(\bar{t}_{\text{par}})}, z.T_P, x_f.\bar{y}.\bar{t}) \rightsquigarrow \mathcal{U}_{[\mathcal{U}_{C(\bar{t}_{\text{par}})}/z]T_P} \\
\\
\text{REDCASTUPDOWN} \\
\langle T \Leftarrow ?_{\text{Type}_i} \rangle \langle ?_{\text{Type}_i} \Leftarrow \text{germ}_i(\mathbf{h}) \rangle t \rightsquigarrow \langle T \Leftarrow \text{germ}_i(\mathbf{h}) \rangle t \\
\\
\text{REDCASTIND} \\
\langle C(\bar{t}'_{\text{par}}) \Leftarrow C(\bar{t}_{\text{par}}) \rangle D^C(\bar{t}_{\text{par}'}, \bar{t}_{\text{arg}}) \rightsquigarrow D^C(\bar{t}'_{\text{par}}, \bar{t}'_{\text{arg}}) \\
\text{where } \bar{t}'_{\text{arg}_i} := \langle \text{Args}_i(C, i, D)[\bar{t}'_{\text{par}}, \bar{t}'_{\text{args}}] \Leftarrow \text{Args}_i(C, i, D)[\bar{t}_{\text{par}}, \bar{t}_{\text{args}}] \rangle \bar{t}_{\text{arg}_i} \\
\\
\text{REDCASTTYPE} \quad \text{REDCASTDOMERR} \quad \text{REDCASTCODOMERR} \\
\langle \text{Type}_i \Leftarrow \text{Type}_i \rangle T \rightsquigarrow T \quad \langle T \Leftarrow \mathcal{U}_{\text{Type}_i} \rangle t \rightsquigarrow \mathcal{U}_T \quad \langle \mathcal{U}_{\text{Type}_i} \Leftarrow T \rangle t \rightsquigarrow \mathcal{U}_{\mathcal{U}_{\text{Type}_i}} \\
\\
\text{REDCASTHEADERR} \quad \text{REDCASTFUN} \\
\frac{\text{head}(T) \neq \text{head}(T')}{\langle T' \Leftarrow T \rangle t \rightsquigarrow \mathcal{U}_T} \quad \langle (x : T'_1) \rightarrow T'_2 \Leftarrow (x : T_1) \rightarrow T_2 \rangle t \\
\rightsquigarrow \lambda y. \langle T'_2 \Leftarrow [T_1 \Leftarrow T'_1]y/x]T_2 \rangle (t \langle T_1 \Leftarrow T'_1 \rangle y) \\
\\
\text{REDCASTFUNGERM} \\
\frac{(x : T_1) \rightarrow T_2 \neq \text{germ}_j(\Pi) \text{ for } j \geq i}{\langle ?_{\text{Type}_i} \Leftarrow (x : T_1) \rightarrow T_2 \rangle t \rightsquigarrow \langle ?_{\text{Type}_i} \Leftarrow ?_{\text{Type}_i} \rightarrow ?_{\text{Type}_i} \rangle \langle ?_{\text{Type}_i} \rightarrow ?_{\text{Type}_i} \Leftarrow (x : T_1) \rightarrow T_2 \rangle t} \\
\\
\text{REDCASTINDGERM} \\
\frac{C(\bar{t}_{\text{par}}) \neq \text{germ}_j(C) \text{ for } j \geq i}{\langle ?_{\text{Type}_i} \Leftarrow C(\bar{t}_{\text{par}}) \rangle t \rightsquigarrow \langle ?_{\text{Type}_i} \Leftarrow C(?_{\text{Params}(C)}) \rangle \langle C(?_{\text{Params}(C)}) \Leftarrow C(\bar{t}_{\text{par}}) \rangle t}
\end{array}$$

Fig. 4. CastCIC: Typing and Reduction rules that extend BCIC

least precise type with a given head in universe i . For function types, the germ is $?_{@i} \rightarrow ?_{@i}$, and Type_i is its own germ. For inductives, the germ is $C(?_{@i})$ where the i 's are the parameters' levels.

3.3 CastCIC: The Cast Calculus

3.3.1 Syntax, Typing and Reductions. Figure 4 presents CastCIC, the cast calculus for GCIC. CastCIC extends BCIC with the unknown term $?$, an error \mathcal{U} , and a cast $\langle T_2 \Leftarrow T_1 \rangle t$ from type T_1 to T_2 . Forms $?_T$ and \mathcal{U}_T are ascribed with their type T , which affect the dynamic semantics of CastCIC. The CastCIC type system contains all the rules from Fig. 2 plus the rules of Fig. 4. Terms $?_T$ and \mathcal{U}_T synthesize their ascribed type T , while casts synthesize the destination type, given that the term being cast checks against the source type, and that both types are well-formed. Because casts are explicit, the checking rule uses definitional equality, rather than consistency.

The CastCIC semantics includes all BCIC reductions, plus cast rules and “propagation rules” that handle $?$ and \mathcal{U} . At type $(x : T_1) \rightarrow T_2$, $?$ and \mathcal{U} expand to $\lambda(x : T_1). ?_{T_2}$ in REDPROPFUNUNK and $\lambda(x : T_1). \mathcal{U}_{T_2}$ in REDPROPFUNERR. In the remaining REDPROP rules, eliminating or casting $?$ or \mathcal{U} produces $?$ or \mathcal{U} . Cast rules either convert between types with the same head, cast to

$\Gamma \vdash t \rightarrow t \Rightarrow \mathbf{T} \quad \Gamma \vdash t \rightarrow t \Leftarrow \mathbf{T} \quad \Gamma \vdash t \rightarrow t \Rightarrow_{\mathbf{h}} \mathbf{T} \quad (\text{Elaboration})$	
$\frac{\text{ELABUNK}}{\Gamma \vdash ?_{@i} \rightarrow ?_{?_{\text{Type}_i}} \Rightarrow ?_{\text{Type}_i}}$ $\frac{\text{ELABUNKFUN} \quad \Gamma \vdash t \rightarrow t \Rightarrow ?_{\text{Type}_i}}{\Gamma \vdash t \rightarrow \langle ?_{\text{Type}_i} \rightarrow ?_{\text{Type}_i} \Leftarrow ?_{\text{Type}_i} \rangle t \Rightarrow_{\Pi} ?_{\text{Type}_i} \rightarrow ?_{\text{Type}_i}}$ $\frac{\text{ELABCST} \quad \Gamma \vdash t \rightarrow t \Rightarrow \mathbf{T}' \quad \mathbf{T}' \cong_{\rightarrow} \mathbf{T}}{\Gamma \vdash t \rightarrow \langle \mathbf{T} \Leftarrow \mathbf{T}' \rangle t \Leftarrow \mathbf{T}}$ $\frac{\text{ELABAPP} \quad \Gamma \vdash t_0 \rightarrow t_0 \Rightarrow_{\Pi} (x : \mathbf{T}_1) \rightarrow \mathbf{T}_2 \quad \Gamma \vdash t_1 \rightarrow t_1 \Leftarrow \mathbf{T}_1}{\Gamma \vdash t_0 t_1 \rightarrow t_0 t_1 \Rightarrow [t_1/x] \mathbf{T}_2}$ $\frac{\text{ELABUNKIND} \quad \Gamma \vdash t \rightarrow t \Rightarrow ?_i}{\Gamma \vdash t \rightarrow \langle \mathbf{C}(?_{\text{Params}(\mathbf{C},i)}) \Leftarrow ?_i \rangle t \Leftarrow \mathbf{C}(\text{Params}(\mathbf{C},i))}$	$\frac{\text{ELABUNKUNIV} \quad \Gamma \vdash t \rightarrow t \Rightarrow ?_{i+1}}{\Gamma \vdash t \rightarrow \langle \text{Type}_i \Leftarrow ?_{i+1} \rangle t \Rightarrow_{\text{Type}} \text{Type}_i}$

Fig. 5. Elaboration from GCIC to CastCIC (homomorphic rules omitted)

$?_{\text{Type}_i}$, or produce an error. A cast from Type_i to itself reduces away (REDCASTTYPE). For inductives, casts from $\mathbf{C}(\bar{t}_1)$ to $\mathbf{C}(\bar{t}_2)$ reduce by casting arguments to their new types (REDCASTIND). Note that t_{par} and $t_{\text{par}'}$ need not match, but typing guarantees that they are convertible. Casts between types with mismatched heads produce an error (REDCASTHEADERR), as do casts to or from $\mathbf{U}_{\text{Type}_i}$ (REDCASTDOMERR, REDCASTCODOMERR). A cast from the germ for a given head does *not* reduce: $\langle ?_{\text{Type}_i} \Leftarrow \text{germ}(\mathbf{T}) \rangle$ acts as a tag, injecting into $?_{\text{Type}_i}$. Casts from non-germ types to $?_{\text{Type}_i}$ decompose into casts through the germ that are then tagged with their injection into $?_{\text{Type}_i}$ (REDCASTFUNGERM, REDCASTINDGERM). In REDCASTUPDOWN, a cast *from* $?_{\text{Type}_i}$ to \mathbf{T} reduces when the value being cast originates from a type with a matching head, and was accordingly tagged with a cast from head(\mathbf{T}) to $?_{\text{Type}_i}$.

3.3.2 Elaboration. Finally, elaboration (Fig. 5) defines the relationship between GCIC and CastCIC. Like CastCIC, elaboration has synthesis, checking, and constrained synthesis, but each produces the elaboration of the subject term as output. ELABUNK synthesizes the unknown type for the unknown term at the given universe level. ELABAPP works like a normal dependent function application, but uses the elaboration of the argument to replace the parameter in the return type. ELABCST checks a term against a type consistent with its synthesized type, inserting the cast between these types into the elaboration. Figure 5 also defines new constrained synthesis rules. Rule ELABUNKFUN works like the corresponding lemma, but adds the necessary cast to the elaborated term. Rules ELABUNKIND and ELABUNKUNIV work similarly. We omit the elaboration rules corresponding to the remaining BCIC rules, which homomorphically elaborate the sub-terms of a given term.

Elaboration defines GCIC typing: we say $t : T$ when $\cdot \vdash T \rightarrow \mathbf{T} \Rightarrow_{\text{Type}} \text{Type}_i$ and $\cdot \vdash t \rightarrow t \Leftarrow \mathbf{T}$.

4 PROPOSITIONAL EQUALITY

The main contribution of our paper is GEq: an extension of GCIC with propositional equality, where the information about an imprecise value accumulates at run time to detect inconsistencies. We define GEq's semantics using a cast calculus CastEq, which extends CastCIC with equality.

The core idea is that a surface-language proof of type $t_1 =_T t_2$ is elaborated into a witness for the consistency of t_1 and t_2 . Much like evidence³ from AGT [Garcia et al. 2016] or the middle-type from threesomes [Siek and Wadler 2010], the consistency witness between terms is a term that is at least as precise as either term. The standard equality proof, $\text{refl}_t : t =_T t$, witnesses that t is

³Evidence is more complex in AGT, since it can witness subtyping. Evidence for plausible equality between types collapses to a single term as precise as the equated terms.

consistent with itself, while the imprecise proof $?_{t_1=T t_2}$ is witnessed by the least precise term that is dynamically consistent with t_1 and t_2 . As a program runs, equality witnesses may take values between these extremes, which may be more precise than the witness for `refl` when t is imprecise.

The technical challenge with adding propositional equality is determining how to combine information represented by the equality witnesses. When casting between types $t_1 =_T t_2$ to $t'_1 =_T t'_2$, both of which may be imprecise, we must transform a witness t_w for $t_1 =_T t_2$ to one for $t'_1 =_T t'_2$, but even though t_w is as precise as t_1 and t_2 , it may not be as precise as t'_1 and t'_2 . So we need a composition operator that can take t'_1 , t'_2 and t_w and produce a term that is as precise as all three. However, to respect static observational equivalences and avoid the problems of §2.4, composition cannot be a syntactic meta-operation. The issue is with neutral terms, i.e., variables, or terms whose reduction is blocked by applying or eliminating a variable. Syntactic composition would require distinct neutral terms to compose to an error, but that would violate extensionality.

Along with composition, we must define a notion of precision that determines valid witnesses of consistency. For the evolution of type information to be monotone, the operator $\&$ should compute a lower bound with respect to this notion of precision. Computing the greatest lower bound prevents premature errors, although the proof that composition is the greatest lower bound is left to future work. With non-dependent gradual types, precision can be syntactically, by adding structural rules to $t \sqsubseteq ?_T$, but structural rules are not flexible enough to handle composition.

The solutions to these two challenges are interdependent. We avoid the issues with syntactic composition by adding it as a separate syntactic construct to `CastEq`, so that composition of neutral terms is itself a neutral term. However, if composition is a construct in `CastEq`, then precision must be defined to accommodate terms that feature composition, so composing two neutral terms produces something that is actually as precise as those two terms. Precision must be defined to respect composition without losing its other important properties, such as transitivity.

This section gives typing and semantics for gradual propositional equality, where proofs of equality are represented by consistency witnesses, but we leave the exact definitions of consistency and precision unspecified. In §5, we describe the properties that consistency and precision should fulfill to ensure that `GEq` satisfies type safety and the gradual guarantees. Finally, §6 instantiates `GEq` with notions of precision and consistency that fulfill our goals while ensuring decidable consistency-checking. We separate our presentation in this way to motivate the choices we make in the design of precision, and to avoid monolithic proofs when developing `GEq`'s metatheory.

We write precision as $\Gamma_1 | \Gamma_2 \vdash t_1 \sqsubseteq_{\rightarrow} t_2$ and consistency as $t_1 \cong_{\rightarrow} t_2$, highlighting the operators in grey to indicate that their definitions are not yet specified. The subscript \rightarrow on \cong_{\rightarrow} indicates that it is *definitional consistency*, whose name is chosen by analogy to definitional equality, since the operands can be reduced before being compared structurally [Martin-Löf 1975]. Precision is *precision modulo conversion*, meaning it is closed under the equivalence relation given by convertibility. Unlike consistency, precision modulo conversion can look backwards in time, relating terms that are the *results* of reducing syntactically-related terms, in addition to relating terms that are syntactically-related after reducing. We discuss the need for this in §4.2. Precision takes two contexts, as its operands must be typed in different contexts.

4.1 GEq Syntax and Typing

Fig. 6 extends `GCIC` to `GEq` by adding the equality type, introduction form, and eliminator. Their types are identical to what is expected in the static setting. Again, because surface typing is defined by elaboration, the given rules are actually admissible lemmas. An equality type $t_1 =_T t_2$ denotes equality between any two values of consistent types (because each endpoint is checked against T). The reflexive proof `reflt` synthesizes type $t =_T t$, so long as t is well typed at type T . The eliminator

$$\boxed{t ::= t =_T t \mid \mathit{refl}_t \mid \mathbf{J}(x.T, t_1, t_2, t_3, t_4)} \quad \frac{\Gamma \vdash t_{eq} \Rightarrow = t_1 =_T t_2 \quad \Gamma, x : T \vdash T_P \Rightarrow_{\text{Type}} \mathbf{Type}_i}{\Gamma \vdash t_1 \Leftarrow T \quad \Gamma \vdash t_2 \Leftarrow T \quad \Gamma \vdash t_{p1} \Leftarrow [t_1/x] T_P} \\
\boxed{h ::= =} \quad \frac{\Gamma \vdash \mathbf{J}(x.T_P, t_1, t_2, t_{p1}, t_{eq}) \Rightarrow [t_2/x] T_P}{\Gamma \vdash T \Rightarrow_{\text{Type}} \mathbf{Type}_i} \\
\frac{\Gamma \vdash t_1 \Leftarrow T \quad \Gamma \vdash t_2 \Leftarrow T \quad \Gamma \vdash T \Rightarrow_{\text{Type}} \mathbf{Type}_i}{\Gamma \vdash t_1 =_T t_2 \Rightarrow \mathbf{Type}_i} \quad \frac{\Gamma \vdash t \Rightarrow T}{\Gamma \vdash \mathit{refl}_t \Rightarrow t =_T t}$$

Fig. 6. GEq: Syntax and Typing Lemmas

$$\boxed{t ::= t_1 =_T t_2 \mid \mathit{refl}(t)_{t_1 \cong t_2} \mid \mathbf{J}(x.T, t_1, t_2, t_3, t_4) \mid t_1 \&_T t_2} \quad \frac{\text{ELABREFL}}{\Gamma \vdash t \rightarrow t \Rightarrow T} \\
\boxed{h ::= \lambda \mid C^D \mid \mathit{refl} \mid =} \quad \frac{\Gamma \vdash \mathit{refl}_t \rightarrow \mathit{refl}(t)_{t \cong t} \Rightarrow t =_T t}{\Gamma \vdash \mathit{refl}_t \rightarrow \mathit{refl}(t)_{t \cong t} \Rightarrow t =_T t} \\
\text{CASTREFL} \quad \frac{\Gamma \vdash t_w \Rightarrow T \quad \Gamma \vdash t_1 \Leftarrow T \quad \Gamma \vdash t_2 \Leftarrow T}{\Gamma \mid \Gamma \vdash t_w \sqsubseteq t_1 \quad \Gamma \mid \Gamma \vdash t_w \sqsubseteq t_2} \\
\frac{\Gamma \mid \Gamma \vdash t_w \sqsubseteq t_1 \quad \Gamma \mid \Gamma \vdash t_w \sqsubseteq t_2}{\Gamma \vdash \mathit{refl}(t_w)_{t_1 \cong t_2} \Rightarrow t_1 =_T t_2} \quad \text{CASTCOMP} \quad \frac{\Gamma \vdash t_1 \Leftarrow T \quad \Gamma \vdash t_2 \Leftarrow T}{\Gamma \vdash t_1 \&_T t_2 \Rightarrow T}$$

Fig. 7. CastEq: Syntax, Key Typing and Elaboration Rules

\mathbf{J} takes a type T_P parameterized over a value of type T ,⁴ along with two values of type T . Then, given a value of type $[t_1/x]T_P$, and a proof t_{eq} that t_1 and t_2 are equal, the elimination has type $[t_2/x]T_P$. That is, if two values are equal, we can take any term whose type refers to the first, and transform it into a term whose type refers to the second.

4.2 CastEq Syntax and Typing

Fig. 7 extends CastCIC to CastEq by adding propositional equality and the gradual composition operator. We extend the syntax for a static head h to include value constructors, not just types, which is useful when defining the semantics of composition. A proof of reflexivity is written as $\mathit{refl}(t_w)_{t_1 \cong t_2}$, where t_1 and t_2 are the equated terms, and t_w is a witness of the (dynamic) consistency of those endpoints. We borrow the notation $(t_w)_{t_1 \cong t_2}$ from Garcia et al. [2016] to indicate that t_w contains information supporting the (dynamic) consistency of t_1 and t_2 . Composition is ascribed with the type of its arguments so that we can ascribe the proper T when the composition of two terms steps to \mathcal{U}_T .

For typing, CASTCOMP synthesizes a composition's ascribed type when both arguments check against that type. In CASTREFL, $\mathit{refl}(t_w)_{t_1 \cong t_2}$ synthesizes $t_1 =_T t_2$ if the witness t_w is as precise as both t_1 and t_2 . In ELABREFL, refl_t is elaborated into $\mathit{refl}(t)_{t \cong t}$, i.e., a term serves as the initial witness that it is equal to itself. If t is imprecise, casts applied to the equality proof may produce more precise witnesses, but the programmer never constructs a witness directly. We omit typing rules for $=_T$ and \mathbf{J} , as they mirror the lemmas in Fig. 6, as do their elaboration rules.

Precision must be closed under convertibility because, as Lennon-Bertrand et al. [2022] note, syntactic precision is not preserved by stepping the less precise term. Since $?_T$ is less precise than x , the less precise term may reduce in a way that is blocked for the other term. So for contextual steps to preserve CASTREFL, the results of stepping related terms must be related.

$$\begin{array}{c}
\text{REDJ} \\
\mathbf{J}(x.\mathbf{T}_P, t_1, t_2, \mathbf{t}_{P1}, \mathbf{refl}(t_w)_{t_1 \cong t_2'}) \rightsquigarrow \langle [t_2/x]\mathbf{T}_P \Leftarrow [t_w/x]\mathbf{T}_P \rangle \langle [t_w/x]\mathbf{T}_P \Leftarrow [t_1/x]\mathbf{T}_P \rangle \mathbf{t}_{P1} \\
\text{REDEQGERM} \\
\frac{t_1 =_{\mathbf{T}} t_2 \neq \text{germ}_\ell(=)}{\langle ?_{\text{Type}_\ell} \Leftarrow t_1 =_{\mathbf{T}} t_2 \rangle t \rightsquigarrow \langle ?_{\text{Type}_\ell} \Leftarrow ?_{\text{Type}_\ell} =_{\text{Type}_\ell} ?_{\text{Type}_\ell} \rangle \langle ?_{\text{Type}_\ell} =_{\text{Type}_\ell} ?_{\text{Type}_\ell} \Leftarrow t_1 =_{\mathbf{T}} t_2 \rangle t} \\
\begin{array}{cc}
\text{PROPEQUNK} & \text{PROPEQ(UNK,ERR)} \\
?_{t_1 =_{\mathbf{T}} t_2} \rightsquigarrow \mathbf{refl}(t_1 \ \&_{\mathbf{T}} \ t_2)_{t_1 \cong t_2} & \mathcal{U}_{t_1 =_{\mathbf{T}} t_2} \rightsquigarrow \mathbf{refl}(\mathcal{U}_{\mathbf{T}})_{t_1 \cong t_2}
\end{array} \\
\text{REDCASTEQ} \\
\langle t'_1 =_{\mathbf{T}'} t'_2 \Leftarrow t_1 =_{\mathbf{T}} t_2 \rangle \mathbf{refl}(t_w)_{t_1 \cong t_2'} \rightsquigarrow \mathbf{refl}(\langle \langle \mathbf{T}' \Leftarrow \mathbf{T} \rangle t_w \ \&_{\mathbf{T}'} \ t'_1 \ \&_{\mathbf{T}'} \ t'_2 \rangle)_{t_1 \cong t_2'}
\end{array}$$

Fig. 8. CastEq: Reductions for Casts and J

4.3 Cast Semantics

A challenge with gradual equality is designing its dynamic semantics. In a fully static language, **refl** always equates identical values, so **J** performs no computation other than pattern matching on the proof of equality. In the presence of type imprecision, **J** must perform casts. We also need reductions for casts between equality types. Figure 8 gives reductions for **J** and casts. The REDJ rule reduces by casting *through* the motive \mathbf{T}_P with x bound to the witness t_w . The typing of equality guarantees that this witness t_w is as precise as either t_1 and t_2 . So $[t_w/x]\mathbf{T}_P$ is like a middle type, since it is more precise than $[t_1/x]\mathbf{T}_P$ (the type of \mathbf{t}_{P1}) and $[t_2/x]\mathbf{T}_P$ (the type of the result).

Why cast through the middle, and not directly from $[t_1/x]\mathbf{T}_P$ to $[t_2/x]\mathbf{T}_P$? As §2.3 showed, the witness tracks constraints as the program runs, and since composition is monotone, its precision only increases. So constraints are remembered, and **J** only succeeds if the equated terms are consistent with all those remembered constraints, allowing the programmer to see when a static constraint has been dynamically violated. Also, the witness ensures that equalities between inconsistent values cannot be used without flagging an error. Without a witness, one could have $?_{2=N,5}$, despite the type being statically uninhabited. Then **J** could use this equality to convert from **Vec Float** (2 mod 3) to **Vec Float** (5 mod 3): the cast would succeed, despite the absurdity of the initial equality. Going through the middle type catches such absurd cases.

For casting **refl** between equality types, the REDCASTEQ rule first casts the witness to the correct type. The typing rule CASTREFL requires the witness to be as precise as the endpoints, but the result of casting the witness might not fulfill this! So the witness is composed with both endpoints, producing a precision-related result. These casts are precisely why we need a composition operator.

The propagation rules PROPEQUNK and PROPEQERR reduce $?$ and \mathcal{U} at equality types to **refl** with the least and most precise witnesses, respectively. REDEQGERM casts an equality proof to $?_{\text{Type}_\ell}$ by casting through the germ type, just like with functions and constructors.

4.4 Semantics of Composition

Figure 9 gives the semantics of composition. Technically, we do not need composition as an operator in CastEq itself, but only for witnesses and cast type ascriptions. However, because dependent types remove the separation between terms and types, witnesses and cast types need dynamic semantics. So for simplicity, we let witnesses and cast-types be any CastEq terms, and add composition to CastEq's semantics, rather than duplicating CastEq's semantics for a witness-specific language.

Several composition rules receive two or three different sets of type ascriptions, with one on the composition itself, and possibly one on each of the composed terms. In such cases, the choice

⁴The full **J** in type theory parameterizes \mathbf{T}_P over the equality proof. Section 7.1.2 shows why this is not needed for GEq.

$\text{UnkVal } \mathbf{T}$ (<i>Types with ? and \mathbf{U} forms</i>)	$\overline{\text{UnkVal } \mathbf{Type}_\ell}$	$\overline{\text{UnkVal } \mathbf{C}(\bar{t})}$	$\overline{\text{UnkVal } ?_{\mathbf{Type}_\ell}}$
$t_1 \rightsquigarrow t_2$ (<i>Composition reductions</i>)	$\frac{\text{REDCOMPUNKL}}{\text{UnkVal } \mathbf{T}} \quad \frac{}{?_{\mathbf{T}} \&_{\mathbf{T}'} t \rightsquigarrow t}$	$\frac{\text{REDCOMPUNKR}}{\text{UnkVal } \mathbf{T}} \quad \frac{}{t \&_{\mathbf{T}'} ?_{\mathbf{T}} \rightsquigarrow t}$	$\frac{\text{REDCOMPERRL}}{\text{UnkVal } \mathbf{T}} \quad \frac{}{\bar{\mathbf{U}}_{\mathbf{T}} \&_{\mathbf{T}'} t \rightsquigarrow \bar{\mathbf{U}}_{\mathbf{T}}}$
$\frac{\text{REDCOMPERRR}}{\text{UnkVal } \mathbf{T}} \quad \frac{}{t \&_{\mathbf{T}'} \bar{\mathbf{U}}_{\mathbf{T}} \rightsquigarrow \bar{\mathbf{U}}_{\mathbf{T}}}$	$\frac{\text{REDCOMPGERM}}{(\langle ?_{\mathbf{Type}_\ell} \Leftarrow \text{germ}(\mathbf{h}) \rangle t_1) \&_{?_{\mathbf{Type}_\ell}} (\langle ?_{\mathbf{Type}_\ell} \Leftarrow \text{germ}(\mathbf{h}) \rangle t_2) \rightsquigarrow \langle ?_{\mathbf{Type}_\ell} \Leftarrow \text{germ}(\mathbf{h}) \rangle (t_1 \&_{\text{germ}(\mathbf{h})} t_2)}$	$\frac{\text{REDCOMPHEADERR}}{\text{head}(t_1) = \mathbf{h}_1 \quad \text{head}(t_2) = \mathbf{h}_2} \quad \frac{}{t_1 \&_{\mathbf{T}} t_2 \rightsquigarrow \bar{\mathbf{U}}_{\mathbf{T}}}$	
$\frac{\text{REDCOMPGERMERR}}{\langle ?_{\mathbf{Type}_\ell} \Leftarrow \text{germ}(\mathbf{h}_1) \rangle t_1 \&_{?_{\mathbf{Type}_\ell}} (\langle ?_{\mathbf{Type}_\ell} \Leftarrow \text{germ}(\mathbf{h}_2) \rangle t_2) \rightsquigarrow \bar{\mathbf{U}}_{?_{\mathbf{Type}_\ell}}$	$\frac{}{\mathbf{h}_1 \neq \mathbf{h}_2}$	$\frac{}{\mathbf{h}_1 \neq \mathbf{h}_2}$	
$\frac{\text{REDCOMPEQ}}{t'_1 := \langle \mathbf{T}_1 \&_{\mathbf{Type}_\ell} \mathbf{T}_2 \Leftarrow \mathbf{T}_1 \rangle t_1 \quad \&_{(\mathbf{T}_1 \&_{\mathbf{Type}_\ell} \mathbf{T}_2)} \quad \langle \mathbf{T}_1 \&_{\mathbf{Type}_\ell} \mathbf{T}_2 \Leftarrow \mathbf{T}_2 \rangle t_2$ $t'_2 := \langle \mathbf{T}_1 \&_{\mathbf{Type}_\ell} \mathbf{T}_2 \Leftarrow \mathbf{T}_1 \rangle t'_1 \quad \&_{(\mathbf{T}_1 \&_{\mathbf{Type}_\ell} \mathbf{T}_2)} \quad \langle \mathbf{T}_1 \&_{\mathbf{Type}_\ell} \mathbf{T}_2 \Leftarrow \mathbf{T}_2 \rangle t'_2$ $\frac{}{(t_1 =_{\mathbf{T}_1} t'_1) \&_{\mathbf{Type}_\ell} (t_2 =_{\mathbf{T}_2} t'_2) \rightsquigarrow t'_1 =_{(\mathbf{T}_1 \&_{\mathbf{Type}_\ell} \mathbf{T}_2)} t'_2}$	$\frac{\text{REDCOMPLAM}}{t_1 \&_{(x:\mathbf{T}_1) \rightarrow \mathbf{T}_2} t_2 \rightsquigarrow \lambda(x:\mathbf{T}_1). ((t_1 x) \&_{\mathbf{T}_2} (t_2 x))}$		
$\frac{\text{REDCOMPREFL}}{\text{refl}(t_1)_{\vdash t_L \cong t_R} \&_{t'_L = t'_R} \text{refl}(t_2)_{\vdash t'_L \cong t'_R} \rightsquigarrow \text{refl}(t_1 \&_{\mathbf{T}} t_2)_{\vdash t_L \cong t_R}}$			
$\frac{\text{REDCOMPIND}}{\text{C}(t_1) \&_{\mathbf{Type}_\ell} \text{C}(t_2) \rightsquigarrow \text{C}(\&(\text{Params}(\mathbf{C}, i), \bar{t}_1, \bar{t}_2))}$	$\frac{\text{REDCOMPCON}}{\mathbf{D}^{\mathbf{C}}(\bar{t}_1, \bar{t}_2) \&_{\mathbf{C}(\bar{v})} \mathbf{D}^{\mathbf{C}}(\bar{t}'_1, \bar{t}'_2) \rightsquigarrow \mathbf{D}^{\mathbf{C}}(\bar{t}, \&(\text{Args}(\mathbf{C}, i, \mathbf{D}_k), \bar{t}_2, \bar{t}'_2))}$		
$\frac{\text{REDCOMPPI}}{(x:\mathbf{T}_1) \rightarrow \mathbf{T}'_1 \&_{\mathbf{Type}_\ell} (x:\mathbf{T}_2) \rightarrow \mathbf{T}'_2 \rightsquigarrow (x:(\mathbf{T}_1 \&_{\mathbf{Type}_\ell} \mathbf{T}_2)) \rightarrow [\langle \mathbf{T}_1 \Leftarrow \mathbf{T}_1 \&_{\mathbf{Type}_\ell} \mathbf{T}_2 \rangle x/x] \mathbf{T}'_1 \&_{\mathbf{Type}_\ell} [\langle \mathbf{T}_2 \Leftarrow \mathbf{T}_1 \&_{\mathbf{Type}_\ell} \mathbf{T}_2 \rangle x/x] \mathbf{T}'_2}$			
$\&(\overline{(x:\mathbf{T})}, \bar{t}_1, \bar{t}_2)$ (<i>Telescope Composition</i>)			
$\&(\cdot, \cdot, \cdot) := \cdot$	$\&(\overline{(x:\mathbf{T}_1)} \cdot \overline{(y:\mathbf{T}_2)}, t_1 \cdot \bar{t}'_1, t_2 \cdot \bar{t}'_2) :=$ $(t_1 \&_{\mathbf{T}_1} t_2) \cdot \&(\overline{(y: [(t_1 \&_{\mathbf{T}_1} t_2)/x] \mathbf{T}_2}), \text{seq}_1, \text{seq}_2)$ <i>where</i> $\text{seq}_1 := \langle [(t_1 \&_{\mathbf{T}_1} t_2)/x] \mathbf{T}_2 \Leftarrow [t_1/x] \mathbf{T}_2 \rangle t'_1$ $\text{seq}_2 := \langle [(t_1 \&_{\mathbf{T}_1} t_2)/x] \mathbf{T}_2 \Leftarrow [t_2/x] \mathbf{T}_2 \rangle t'_2$		

Fig. 9. CastEq: Semantics for $\&$

of ascription in the reduct is arbitrary. For example, REDCOMPERRL uses the ascription from $\bar{\mathbf{U}}$, rather than what was on $\&$. While we do not require that the ascriptions be syntactically equal, the typing rules ensure that syntactically distinct ascriptions are definitionally equal, so the choice of ascription in the reduct does not affect the final result of evaluation.

Each composition rule resembles a unification rule, where each use of $?$ is treated as a unification variable. Terms that have different head tags compose to an error, and terms that have the same head tag compose by composing their parts. Any place where one term has $?$ but the other term contains more precise information, the precise information is retained in the output. This ensures

that when $t_1 \&_T t_2$ reduces, the result is a term that is as precise as both t_1 and t_2 . For $?_T \&_T t$, we produce t ($\text{REDCOMPUNK}(L,R)$), since t is always as precise as itself and $?_T$. Likewise, the rules that produce \bar{U} satisfy this, since it is the most precise term. We see this in $\text{REDCOMPUNK}(L,R)$, which composes with \bar{U} , and in REDCOMPHEADERR and REDCOMPGERMERR , where composing non-neutral terms with distinct heads reduces to \bar{U} .

The remaining rules compose terms with the same head h , such as when both are functions or both are built with the same D^C . In these cases, the head h is applied to the respective composition of the arguments, e.g., the composition of functions is a function returning the composition of the bodies. For equality proofs and inhabitants of $?_{\text{Type}_e}$, the head can be applied directly (REDCOMPLAM , REDCOMPREFL , REDCOMPGERM). In REDCOMPLAM , composing functions yields a new function that composes the results of the given functions for each argument. In the remaining cases, we must account for how types of later arguments depend on the values of earlier arguments. REDCOMPPI produces a domain by composing the argument domains, which is the type of the parameter x . The codomain x 's each have their own domain types, so we cast all uses of x from the composed type to the expected type. REDCOMPSEQ composes equality types: the type ascriptions are composed, then the equated terms are cast to this composed type. Composing the results of these casts yields the result endpoints.

The most complex rules are REDCOMPIND and REDCOMPCON . Because type and data constructors have dependent function types, their arguments are *telescopes*: the type of later arguments may depend on the *values* of previous parameters and arguments. To compose type or data constructor applications, we compose the parameters and arguments element-wise, but composing two arguments changes the type of later arguments. To compose telescopes, the metafunction $\bar{\&}$ traverses the types of type and data constructors, composing arguments element-wise and adding casts to the bound variables in later arguments.

To see why composing needs casts, consider dependent pairs formulated as inductive types.

$$\begin{aligned} \text{data DPair} & : (X : \text{Type}) \rightarrow (P : X \rightarrow \text{Type}) \rightarrow \text{Type} \text{ where} \\ \text{mkDPair} & : (x : X) \rightarrow P\ x \rightarrow \text{DPair}\ X\ P \end{aligned}$$

One example of a dependent pair type is $\text{DPair} (\mathbb{N} \times \mathbb{N}) (\lambda x. (\pi_1\ x) + (\pi_2\ x) =_{\mathbb{N}} 3)$, i.e., the Curry-Howard equivalent of “there exists a pair of numbers such that adding them yields 3.” Suppose we want to compose two inhabitants of this type, say $\text{mkDPair} (1, ?_{\mathbb{N}}) \text{refl}(3)_{\vdash-1+?_{\mathbb{N}} \cong 3}$ and $\text{mkDPair} (?_{\mathbb{N}}, 2) \text{refl}(3)_{\vdash-?_{\mathbb{N}}+2 \cong 3}$. To compose the first element, we can produce $(1 \&_{\mathbb{N}} ?_{\mathbb{N}}, ?_{\mathbb{N}} \&_{\mathbb{N}} 2)$, which reduces to $(1, 2)$. However, for the second element, the two proofs $\text{refl}(3)_{\vdash-1+?_{\mathbb{N}} \cong 3}$ and $\text{refl}(3)_{\vdash-?_{\mathbb{N}}+2 \cong 3}$ do not have the same type: they equate different terms, so we cannot compose them! Instead, we must first cast each to type $(1 \&_{\mathbb{N}} ?_{\mathbb{N}}) + (?_{\mathbb{N}} \&_{\mathbb{N}} 2) =_{\mathbb{N}} 3$, i.e., the value obtained by replacing x with the composition of the pairs’ first elements in the term $((\pi_1\ x) + (\pi_2\ x) =_{\mathbb{N}} 3)$. This gives a final result of $\text{mkDPair} (1, 2) (\text{refl}(3)_{\vdash-3 \cong 3})$.

5 PARAMETERIZED METATHEORY: CRITERIA FOR PRECISION AND CONSISTENCY

GEq is now defined except for CastEq 's precision and consistency relations. For non-dependent languages, the semantics of precision can be justified either in terms of sets of static terms [Garcia et al. 2016] or in terms of semantic precision [New and Ahmed 2018]. Such justifications are difficult with dependent types. Our approach is different: we define the important metatheoretic criteria for GEq without referring to precision and consistency, then describe the criteria precision and consistency must fulfill to prove the desired metatheoretic properties. Meeting these criteria guides and justifies our definition of precision and consistency (§6). We see precision and consistency as a means to the end of the desired metatheory.

5.1 Stating the Gradual Guarantees

The gradual guarantees state that reducing the precision of a surface term introduces no new static or dynamic errors. However, to state them formally, we must define what precision means for surface terms. We follow [Lennon-Bertrand et al. \[2022\]](#) and define surface precision as the relation generated by $t \sqsubseteq_{\text{Surf}} ?_{@i}$, plus all the usual structural rules. Essentially, $t \sqsubseteq_{\text{Surf}} t'$ holds if we can obtain t' by replacing some parts of t with some $?_{@i}$. To guarantee preservation of typing, we also need such replacements to be *universe adequate* [[Lennon-Bertrand et al. 2022](#)]. We say that the judgment $t \sqsubseteq_{\text{Surf}} t'$ is universe adequate if, for every subterm r of t , when $\Gamma \vdash r \rightarrow \mathbf{r} \Leftarrow \mathbf{T}$ and $\Gamma \vdash \mathbf{T} \Rightarrow_{\text{Type}} \text{Type}_i$, then any uses of $r \sqsubseteq_{\text{Surf}} ?_{@j}$ have $i = j$. This essentially says t is t' with some subterms replaced by $?_{@i}$ for the right i . We can now state the static gradual guarantee:

Definition 5.1 (Static Gradual Guarantee). If $\cdot \vdash t : T$ and $t \sqsubseteq_{\text{Surf}} t'$ universe-adequately, then $\cdot \vdash t' : T$.

That is, reducing the precision of a program causes no new type errors.

To state the dynamic guarantee without referring to CastEq precision, we must formalize what it means to introduce no new dynamic errors. We follow [New and Ahmed \[2018\]](#) and do this with *semantic precision*, which compares terms by quantifying over all possible boolean contexts. We use booleans because of their simplicity: gradual booleans have only the four values **true**, **false**, $?_{\mathbb{B}}$ and $\mathbb{U}_{\mathbb{B}}$. If a context exists such that reducing a term's precision changes the result from **true** to **false**, then we have violated the guarantee that precision only affects behavior via errors. Likewise, if a context exists such that reducing precision turns **true** to $\mathbb{U}_{\mathbb{B}}$, then reducing precision introduced a new error. By defining semantic precision in terms of *all* contexts, we capture the idea that the above behaviors are impossible for precision-related terms. We formalize this as follows:

Definition 5.2 (Semantic Precision). Boolean precision $\sqsubseteq_{\mathbb{B}}$ is defined by **true** $\sqsubseteq_{\mathbb{B}}$ **true**, **false** $\sqsubseteq_{\mathbb{B}}$ **false**, $\mathbb{U}_{\mathbb{B}}$ $\sqsubseteq_{\mathbb{B}}$ **b**, and **b** $\sqsubseteq_{\mathbb{B}}$ $?_{\mathbb{B}}$ for all $\mathbf{b} : \mathbb{B}$. Then two closed terms are related by *semantic precision*, written $\vDash t \sqsubseteq^{\#} t' : T$ if, for all $C : T \rightarrow \mathbb{B}$, whenever $C[t] \rightarrow^* \mathbf{b}$, then $C[t'] \rightarrow^* \mathbf{b}'$ and $\mathbf{b} \sqsubseteq_{\mathbb{B}} \mathbf{b}'$.

Then the gradual guarantee states that reducing a surface term's precision causes a corresponding reduction in the semantic precision of the surface terms' elaborations.

Definition 5.3 (Dynamic Gradual Guarantee). Suppose $\cdot \vdash t \rightarrow t \Leftarrow T$ and $\cdot \vdash t' \rightarrow t' \Leftarrow T$. If $t \sqsubseteq_{\text{Surf}} t'$ universe-adequately, then $\Gamma \vDash t \sqsubseteq^{\#} t'$.

5.2 Necessary Properties of Precision and Consistency

Next, we list properties that, if satisfied by $\sqsubseteq_{\rightarrow}^{\leftarrow}$ and \cong_{\rightarrow} , suffice to prove type safety, conservative extension of CIC, well-typedness of elaboration, and the gradual guarantees. Each criterion is accompanied by a specific case of the safety or gradual guarantee proofs that motivates its inclusion. We also include criteria that composition should satisfy. While the semantics appear in a prior section (§4.4), the criteria are new from GCIC, so we list them to highlight our contribution.

Though $\sqsubseteq_{\rightarrow}^{\leftarrow}$ is ideal for typing witnesses, it is too lenient to express the monotonicity properties of CastEq. In particular, if the monotonicity of reduction is phrased with $\sqsubseteq_{\rightarrow}^{\leftarrow}$, then consistency must also be closed under convertibility, which would make it undecidable even when its operands terminate. So we introduce a strictly stronger relation $\sqsubseteq_{\rightarrow}$, which, like \cong_{\rightarrow} , only compares after reductions, and not before. This distinction is acceptable because the precision side-condition of **CASTREFL** is never used to prove safety or monotonicity. Rather, the side-condition ensures that witnesses always entail at least as much information as the equated terms. Including the side-condition in CastEq's type conveniently captures the invariant that, when refl_t is elaborated with initial witness t , future witnesses are never lose the information from t .

Safety and Elaboration. For elaboration to preserve types, precision must be reflexive so the initial witness for *refl* is valid. For safety, progress requires that each well-typed non-value can step. So composition must step for each non-value. For preservation, each reduction must preserve types, including composition reductions. If composition yields a precision lower-bound and precision is transitive, the side-conditions of *CASTREFL* can be preserved.

LEMMA 5.4 (PRECISION REFLEXIVE). *If $\Gamma \vdash t \Leftarrow T$ then $\Gamma|\Gamma \vdash t \sqsubseteq_{\rightarrow} t$ (For *ELABREFL* to produce an elaboration that satisfies the $\sqsubseteq_{\rightarrow}$ side-condition of *CASTREFL*).*

LEMMA 5.5 (COMPOSITION SAFETY). *If $t_1 \&_T t_2$ is not a value and $\Gamma \vdash t_1 \&_T t_2 \Leftarrow T$, then $t_1 \&_T t_2 \rightarrow t_3$ for some t_3 and $\Gamma \vdash t_3 \Leftarrow T$ (For progress and preservation)*

LEMMA 5.6 (COMPOSITION CONFLUENCE). *If $t_1 \&_T t_2 \Rightarrow t_3$ and $t_1 \&_T t_2 \Rightarrow t'_3$ maximally, then $t_3 \Rightarrow t'_3$, where \Rightarrow is the parallel reduction relation, standard in confluence proofs [Takahashi 1995] (For confluence, which is needed to show that β -reductions preserve types);*

LEMMA 5.7 (COMPOSITION LOWER BOUND). *If $\Gamma \vdash t_1 \&_T t_2 \Leftarrow T$, then $\Gamma|\Gamma \vdash t_1 \&_T t_2 \sqsubseteq_{\rightarrow} t_1$ and $\Gamma|\Gamma \vdash t_1 \&_T t_2 \sqsubseteq_{\rightarrow} t_2$ (Preserving the $\sqsubseteq_{\rightarrow}$ condition of *CASTREFL* for reduction *REDCASTEQ*);*

LEMMA 5.8 (PRECISION TRANSITIVE). *If $\Gamma_1|\Gamma_2 \vdash t_1 \sqsubseteq_{\rightarrow} t_2$ and $\Gamma_2|\Gamma_3 \vdash t_2 \sqsubseteq_{\rightarrow} t_3 : T$ then $\Gamma_1|\Gamma_3 \vdash t_1 \sqsubseteq_{\rightarrow} t_3$ (Preserving the $\sqsubseteq_{\rightarrow}$ side-condition of *CASTREFL* for reduction *REDCASTEQ*);*

LEMMA 5.9 (PRECISION MODULO CONVERSION). *If $\Gamma_1|\Gamma_2 \vdash t_1 \sqsubseteq_{\rightarrow} t_2$, where $t_2 \rightarrow^* t'_2$, then $\Gamma_1|\Gamma_2 \vdash t_1 \sqsubseteq_{\rightarrow} t'_2$ (Preservation of *CASTREFL* under contextual reduction)*

Conservativity. If *GEq* is to conservatively extend *CIC*, then a fully-static program should be well-typed in *CIC* if and only if it is well-typed in *GEq*. For the most part, the rules only differ when $?$ is involved, but the major exception is *ELABCST*, which let us replace a type with any consistent type (after conversion). So for fully static terms, consistency should coincide with syntactic equality.

LEMMA 5.10 (STATIC CONSISTENCY). *For any static terms t_1 and t_2 , let t_1 and t_2 be their embedding in *CastEq*. Then $t_1 \cong_{\rightarrow} t_2$ iff $t_1 =_{\alpha\beta} t_2$, i.e., if they are statically definitionally equal (For *GEq* to conservatively extend *CIC*).*

Monotonicity. The last group of properties relate to the gradual guarantees. The dynamic gradual guarantee requires that evaluating precision-related terms produces precision-related results. Because of the dependency in dependent types, proving the static guarantee relies on the proof of the dynamic guarantee: *ELABCST* reduces types before comparing for consistency, so precision of types before reduction should be preserved, and reducing the precision of a type should make it consistent with no fewer types. Likewise, to show the static guarantee, elaboration must be monotone in both synthesized types and elaborated terms, since dependent application uses the argument's elaboration in the return type.

LEMMA 5.11 (CAST MONOTONICITY). *Suppose that $\Gamma_1|\Gamma_2 \vdash t_1 \sqsubseteq_{\rightarrow} t_2$, $\Gamma_1 \vdash t_1 \Rightarrow T_1$ and $\Gamma_2 \vdash t_2 \Rightarrow T_2$ where $\Gamma_1|\Gamma_1 \vdash T_1 \sqsubseteq_{\rightarrow} T'_1$ and $\Gamma_2|\Gamma_2 \vdash T_2 \sqsubseteq_{\rightarrow} T'_2$. Then $\Gamma_1|\Gamma_2 \vdash \langle T'_1 \Leftarrow T_1 \rangle t_1 \sqsubseteq_{\rightarrow} \langle T'_2 \Leftarrow T_2 \rangle t_2$ (For *ELABCST* to produce $\sqsubseteq_{\rightarrow}$ -related elaborations for $\sqsubseteq_{\text{surf}}$ -related inputs)*

LEMMA 5.12 (SUBSTITUTION MONOTONE). *Suppose $\Gamma_1|\Gamma_2 \vdash t_1 \sqsubseteq_{\rightarrow} t_2$, where $\Gamma_1 \vdash t_1 \Rightarrow T_1$ and $\Gamma_2 \vdash t_2 \Rightarrow T_2$. If $\Gamma_1(x : T_1)\Delta_1|\Gamma_2(x : T_2)\Delta_2 \vdash t'_1 \sqsubseteq_{\rightarrow} t'_2$, then $\Gamma_1[t_1/x]\Delta_1|\Gamma_2[t_2/x]\Delta_2 \vdash [t_1/x]t'_1 \sqsubseteq_{\rightarrow} [t_2/x]t'_2$ (For *ELABAPP* to be monotone in the return type)*

LEMMA 5.13 (REDUCTION MONOTONE). *If $\Gamma_1|\Gamma_2 \vdash t_1 \sqsubseteq_{\rightarrow} t_2$ and $t_1 \rightarrow^* t'_1$, then $t_2 \rightarrow^* t'_2$ for some t'_2 where $\Gamma_1|\Gamma_2 \vdash t'_1 \sqsubseteq_{\rightarrow} t'_2$ (For *DGG*, to preserve *ELABCST* when reducing precision, and to preserve typing under contextual reduction of $\text{refl}(t_w)_{t_1 \cong t_2}$)*

LEMMA 5.14 (CONSISTENCY MONOTONE FOR PRECISION). *If $\Gamma \vdash t_1 \sqsubseteq_{\rightarrow} t'_1$ and $\Gamma \vdash t_2 \sqsubseteq_{\rightarrow} t'_2$, and $t_1 \cong_{\rightarrow} t_2$, then $t'_1 \cong_{\rightarrow} t'_2$ (So reducing precision of \mathbf{V} and \mathbf{V}' preserves ELABCST)*

LEMMA 5.15 (STRUCTURAL PRECISION). *$\sqsubseteq_{\rightarrow}$ contains all structural rules (For homomorphic elaboration rules to produce $\sqsubseteq_{\rightarrow}$ -related elaboration for $\sqsubseteq_{\text{Surf}}$ -related inputs)*

5.3 Metatheory: Proving Safety and the Gradual Guarantees

Finally, we summarize the properties that we can prove by assuming GEq satisfies the criteria of §5.2. The general idea is that each case in the proofs either (1) is the same as the proof for GCIC [Lennon-Bertrand et al. 2022] or (2) follows directly from one of our criteria. Full proofs can be found in the appendix of the extended technical report [Eremondi et al. 2022].

5.3.1 *Type Safety.* Type safety is shown in the usual way for operational semantics, via progress and preservation [Wright and Felleisen 1994]. Each well-typed CastEq term is either a value, or can step to a well-typed term. Confluence is necessary to prove preservation for dependent types. Space restrictions mean that the formalization of values \mathbf{v} in GEq are in the appendix of the extended technical report [Eremondi et al. 2022], but the idea is to follow Lennon-Bertrand et al. [2022], adding $t_1 \&_{\mathbf{T}} t_2$ as value when t_1 and t_2 are values, neither of t_1 and t_2 is $?$ or \mathbf{U} , and \mathbf{T} is not a function type.

PROPOSITION 5.16 (CONFLUENCE, PROGRESS, PRESERVATION AND ELABORATION).

- \rightarrow is confluent.
- If $\Gamma \vdash t \Leftarrow \mathbf{T}$, then t is a value or $t \rightarrow t'$ for some t' .
- If $\Gamma \vdash t_1 \Leftarrow \mathbf{T}$ and $t_1 \rightarrow t_2$ then $\Gamma \vdash t_2 \Leftarrow \mathbf{T}$.
- If $\Gamma \vdash t \rightarrow t \Leftarrow \mathbf{T}$, then $\Gamma \vdash t \Leftarrow \mathbf{T}$.

These together yield the main safety theorem.

THEOREM 5.17 (TYPE SAFETY). *If $\cdot \vdash t : T$, then t has an elaboration that either steps to a normal form or steps indefinitely.*

As a corollary, we can perform inversion on the typing derivations to obtain weak canonicity. That is, every well-typed closed term that terminates steps to a canonical term of its type.

COROLLARY 5.18 (WEAK CANONICITY). *Suppose $\cdot \vdash t : \mathbf{V}$. Then either t diverges, or $t \rightarrow^* \mathbf{v}$ where \mathbf{v} is $?_{\mathbf{V}}$ or $\mathbf{U}_{\mathbf{V}}$, or the following hold:*

- If \mathbf{V} is $(x : T_1) \rightarrow T_2$ then \mathbf{v} is $\lambda x. t'$
- If \mathbf{V} is $\mathbf{C}_{@[i]}(\overline{t_1})$ then \mathbf{v} is $\mathbf{D}^{\mathbf{C}}_{@[i]}(\overline{t_2})$ for some \mathbf{D} .
- If \mathbf{V} is $t_1 =_{\mathbf{T}} t_2$ then \mathbf{v} is $\text{refl}(t')_{t_1 \cong t_2}$
- If \mathbf{V} is Type_i then \mathbf{v} is one of $\mathbf{C}_{@[i]}(\overline{t_1})$, $(x : T_1) \rightarrow T_2$, Type_{i-1} or $t_1 =_{\mathbf{T}} t_2$.

5.3.2 *Conservatively Extending CIC.* Each CIC rule has a direct analogue in CastEq, so it is clear that it extends CIC. Since most of the gradual-specific rules refer to $?$ or \mathbf{U} , knowing that consistency collapses to α -equivalence on static terms is enough to show that said extension is conservative.

THEOREM 5.19 (CONSERVATIVITY). *For any BCIC-terms t and T , let t and T be the GEq terms corresponding to t and T by mapping BCIC λ to GEq λ , etc. Then $\cdot \vdash t \Leftarrow T$ iff $\cdot \vdash t : T$.*

5.3.3 *Gradual Guarantees.* To prove the gradual guarantees, we use the gradual criteria to show that elaboration is monotone. This, when combined with the monotonicity of \rightsquigarrow with respect to semantic precision, gives us both the static and dynamic guarantees.

PROPOSITION 5.20 (ELABORATION GRADUAL GUARANTEE). *Suppose $t_1 \sqsubseteq_{\text{Surf}} t_2$ and $\Gamma_1 \sqsubseteq_{\rightarrow} \Gamma_2$ (i.e. entries in Γ_1 and Γ_2 are respectively related by $\sqsubseteq_{\rightarrow}$). Then:*

- If $\Gamma_1 \vdash t_1 \rightarrow t_1 \Leftarrow T$ then $\Gamma_2 \vdash t_2 \rightarrow t_2 \Leftarrow T$ for some t_2 where $\Gamma_1 | \Gamma_2 \vdash t_1 \sqsubseteq_{\rightarrow} t_2$.
- If $\Gamma_1 \vdash t_1 \rightarrow t_1 \Rightarrow T_1$ then $\Gamma_2 \vdash t_2 \rightarrow t_2 \Rightarrow T_2$ for some T_2, t_2 where $\Gamma_2 \vdash T_2 \Rightarrow_{\text{Type}} \text{Type}_\ell$ and $\Gamma_1 | \Gamma_2 \vdash t_1 \sqsubseteq_{\rightarrow} t_2$.

This, combined with the preservation of precision under evaluation, is enough to prove the static and dynamic gradual guarantees as stated in §5.1. The hard work lies in proving that reduction preserves precision, which we leave to §6.3.

6 CONSISTENCY AND PRECISION

Motivated by the criteria of §5.2, in this section we extend GCIC's precision and consistency relations to accommodate propositional equality and composition. We show that our relations fulfill the laws of §5.2, thus showing that GEq fulfills type safety and the gradual guarantees, justifying the design of precision and consistency.

6.1 Review: Precision and Consistency in GCIC

6.1.1 *Structural Precision.* Figure 10 recalls *structural precision* from GCIC [Lennon-Bertrand et al. 2022], written as $\Gamma \vdash t_1 \sqsubseteq_{\alpha} t_2$. Structural precision is the syntactic relation out of which definitional precision $\sqsubseteq_{\rightarrow}$ is built. The *generating rules* GENUNK and GENERR establish $?_{\mathbf{T}}$ and $\mathbf{U}_{\mathbf{T}}$ as the least and most precise terms of type \mathbf{T} . For technical reasons, GENUNKUNIV allows some cumulativity for $?_{\text{Type}_\ell}$, while GENERRLAM encodes a version of η -expansion for errors. The *diagonal rules* (named DIAG^*) are structural: terms are precision related if they are built with the same syntactic construct and the corresponding sub-terms are precision-related. We show a few examples, but omit most diagonal rules for space reasons. Finally, *cast rules* capture non-structural properties of casts. Rule CAST-L states that a casting t is more precise than t' if the cast's source and destination types are both more precise than the type of t' , and if t is more precise than t' . The rule CAST-R says the opposite: casting t is less precise than t if the source and destination are both less precise than the type of t and t itself is less precise than t' .

Structural precision uses an auxiliary type judgment: *presynthesis* $\Gamma \vdash t \Rightarrow^* \mathbf{T}$ is defined to be exactly the type synthesis relation without the $\sqsubseteq_{\rightarrow}$ side-condition in CASTREFL. Presynthesis types strictly more terms than synthesis, and both produce the same type, since they differ only in side-conditions. The side-condition is not used in the type-safety proof, so any run-time terms that presynthesize a type are safe. Unlike GCIC, GEq uses precision to type equality witnesses, so presynthesis avoids a circular dependency between typing and precision.

Structural precision is defined mutually with definitional precision (Fig. 10) $\sqsubseteq_{\rightarrow}$, which acts as $\sqsubseteq_{\rightarrow}$ from §5. Definitional precision allows reducing before comparing, and is used with type ascriptions, such as for functions, equality proofs and casts. Since the checking rule for CastCIC allowed arbitrary reductions, a term may be well-typed even if its type ascriptions are not fully reduced. Type ascriptions on a term may need to be reduced before structural precision is apparent. This definition is due to Lennon-Bertrand et al. [2022].

6.1.2 *Syntactic Consistency.* Figure 10 defines consistency for GCIC. All terms are consistent with $?_{\mathbf{T}}$ (CSTUNKL, CSTUNKR). Each syntactic construct also has an (omitted) structural rule. Unlike precision, consistency between terms ignores type ascriptions, and casts are also ignored (CSTCASTL, CSTCASTR). We follow GCIC and let $\mathbf{U}_{\mathbf{T}}$ be consistent with $?_{\mathbf{T}}$.

$\Gamma \vdash \mathbf{t}_1 \sqsubseteq_{\alpha} \mathbf{t}_2$ (Precision: Key Rules from GCIC)				
$\frac{\text{GENUNK} \quad \Gamma_1 \vdash \mathbf{t} \Rightarrow^* \mathbf{T}' \quad \Gamma_1 \Gamma_2 \vdash \mathbf{T}' \sqsubseteq_{\rightarrow} \mathbf{T}}{\Gamma_1 \Gamma_2 \vdash \mathbf{t} \sqsubseteq_{\alpha} ?_{\mathbf{T}}}$	$\frac{\text{GENERR} \quad \Gamma_1 \vdash \mathbf{t} \Rightarrow^* \mathbf{T}' \quad \Gamma_1 \Gamma_2 \vdash \mathbf{T} \sqsubseteq_{\rightarrow} \mathbf{T}'}{\Gamma_1 \Gamma_2 \vdash \mathbf{U}_{\mathbf{T}} \sqsubseteq_{\alpha} \mathbf{t}}$	$\frac{\text{GENUNKUNIV} \quad \Gamma_1 \vdash \mathbf{T} \Rightarrow^*_{\text{Type}} \text{Type}_i \quad i < j}{\Gamma_1 \Gamma_2 \vdash \mathbf{T} \sqsubseteq_{\alpha} ?_{\text{Type}_j}}$	$\frac{\text{DIAGVAR}}{\Gamma_1 \Gamma_2 \vdash \mathbf{x} \sqsubseteq_{\alpha} \mathbf{x}}$	
$\frac{\text{GENERRLAM} \quad \Gamma_2 \vdash \mathbf{t}_2 \Rightarrow^*_{\Pi} (\mathbf{x} : \mathbf{T}_2) \rightarrow \mathbf{T}'_2 \quad \Gamma_1 \Gamma_2 \vdash (\mathbf{x} : \mathbf{T}_1) \rightarrow \mathbf{T}'_1 \sqsubseteq_{\rightarrow} (\mathbf{x} : \mathbf{T}_2) \rightarrow \mathbf{T}'_2}{\Gamma_1 \Gamma_2 \vdash \lambda(\mathbf{x} : \mathbf{T}_1). \mathbf{U}_{\mathbf{T}'_1} \sqsubseteq_{\alpha} \mathbf{t}_2}$		$\frac{\text{DIAGABS} \quad \Gamma_1 \Gamma_2 \vdash \mathbf{T}_1 \sqsubseteq_{\rightarrow} \mathbf{T}_2 \quad \Gamma_1, (\mathbf{x} : \mathbf{T}_1) \Gamma_2, (\mathbf{x} : \mathbf{T}_2) \vdash \mathbf{t}_1 \sqsubseteq_{\alpha} \mathbf{t}_2}{\Gamma_1 \Gamma_2 \vdash \lambda(\mathbf{x} : \mathbf{T}_1). \mathbf{t}_1 \sqsubseteq_{\alpha} \lambda(\mathbf{x} : \mathbf{T}_2). \mathbf{t}_2}$		
$\frac{\text{DIAGCAST} \quad \Gamma_1 \Gamma_2 \vdash \mathbf{T}_1 \sqsubseteq_{\alpha} \mathbf{T}_2 \quad \Gamma_1 \Gamma_2 \vdash \mathbf{T}'_1 \sqsubseteq_{\alpha} \mathbf{T}'_2 \quad \Gamma_1 \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\alpha} \mathbf{t}_2}{\Gamma_1 \Gamma_2 \vdash \langle \mathbf{T}'_1 \Leftarrow \mathbf{T}_1 \rangle \mathbf{t}_1 \sqsubseteq_{\alpha} \langle \mathbf{T}'_2 \Leftarrow \mathbf{T}_2 \rangle \mathbf{t}_2}$				
$\frac{\text{CASTL} \quad \Gamma_2 \vdash \mathbf{t}_2 \Rightarrow^* \mathbf{T}_2 \quad \Gamma_1 \Gamma_2 \vdash \mathbf{T}_1 \sqsubseteq_{\rightarrow} \mathbf{T}_2 \quad \Gamma_1 \Gamma_2 \vdash \mathbf{T}'_1 \sqsubseteq_{\rightarrow} \mathbf{T}_2 \quad \Gamma_1 \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\alpha} \mathbf{t}_2}{\Gamma_1 \Gamma_2 \vdash \langle \mathbf{T}'_1 \Leftarrow \mathbf{T}_1 \rangle \mathbf{t}_1 \sqsubseteq_{\alpha} \mathbf{t}_2}$		$\frac{\text{CASTR} \quad \Gamma_1 \vdash \mathbf{t}_1 \Rightarrow^* \mathbf{T}_1 \quad \Gamma_1 \Gamma_2 \vdash \mathbf{T}_1 \sqsubseteq_{\rightarrow} \mathbf{T}_2 \quad \Gamma_1 \Gamma_2 \vdash \mathbf{T}_1 \sqsubseteq_{\rightarrow} \mathbf{T}'_2 \quad \Gamma_1 \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\alpha} \mathbf{t}_2}{\Gamma_1 \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\alpha} \langle \mathbf{T}'_2 \Leftarrow \mathbf{T}_2 \rangle \mathbf{t}_2}$		
$\Gamma_1 \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\rightarrow} \mathbf{t}_2$ (Definitional Precision)				
$\frac{\text{DEFBASE} \quad \Gamma_1 \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\alpha} \mathbf{t}_2}{\Gamma_1 \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\rightarrow} \mathbf{t}_2}$	$\frac{\text{DEFSTEPL} \quad \mathbf{t}_1 \longrightarrow \mathbf{t}'_1 \quad \Gamma_1 \Gamma_2 \vdash \mathbf{t}'_1 \sqsubseteq_{\rightarrow} \mathbf{t}_2}{\Gamma_1 \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\rightarrow} \mathbf{t}_2}$	$\frac{\text{DEFSTEPR} \quad \mathbf{t}_2 \longrightarrow \mathbf{t}'_2 \quad \Gamma_1 \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\rightarrow} \mathbf{t}'_2}{\Gamma_1 \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\rightarrow} \mathbf{t}_2}$		
$\mathbf{t}_1 \cong_{\alpha} \mathbf{t}_2$ (Consistency: Non-structural Rules)				
$\frac{\text{CSTVAR}}{\mathbf{x} \cong_{\alpha} \mathbf{x}}$	$\frac{\text{CSTUNKL}}{?_{\mathbf{T}} \cong_{\alpha} \mathbf{T}}$	$\frac{\text{CSTUNKR}}{\mathbf{t} \cong_{\alpha} ?_{\mathbf{T}}}$	$\frac{\text{CSTCASTL} \quad \mathbf{t} \cong_{\alpha} \mathbf{t}'}{\langle \mathbf{T}_2 \Leftarrow \mathbf{T}_1 \rangle \mathbf{t} \cong_{\alpha} \mathbf{t}'}$	$\frac{\text{CSTCASTR} \quad \mathbf{t} \cong_{\alpha} \mathbf{t}'}{\mathbf{t} \cong_{\alpha} \langle \mathbf{T}_2 \Leftarrow \mathbf{T}_1 \rangle \mathbf{t}'}$

Fig. 10. Structural Precision and Consistency for GCIC: Key Rules

6.2 Precision and Consistency for GEq

The structural precision laws are not sufficient for handling composition. In particular, we want $\Gamma | \Gamma' \vdash \mathbf{t}_1 \&_{\mathbf{T}} \mathbf{t}_2 \sqsubseteq \mathbf{t}_1$, with the same holding for \mathbf{t}_2 . However, this fact is not derivable from the diagonal rule for composition. Instead, we must add rules to ensure that composing produces a lower bound. However, once we start adding non-structural rules, we must be careful not to disrupt the other properties we need from precision. For example, §5.2 states that precision must be transitive. If we have $(\mathbf{x} \&_{\mathbf{T}} \mathbf{y}) \&_{\mathbf{T}} \mathbf{z} \sqsubseteq \mathbf{x} \&_{\mathbf{T}} \mathbf{y}$ and $\mathbf{x} \&_{\mathbf{T}} \mathbf{y} \sqsubseteq \mathbf{x}$, but we also want $(\mathbf{x} \&_{\mathbf{T}} \mathbf{y}) \&_{\mathbf{T}} \mathbf{z} \sqsubseteq \mathbf{x}$, then we must transitively apply the fact that composing produces a lower bound.

Figure 11 shows the added rules. `DIAGREFL` and `DIAGCOMP`, along with the omitted `DIAGEQ` and `DIAGJ`, are like the other diagonal rules. The rules `PREC COMPL` and `PREC COMP R` encode that the composition is a precision-lower bound, but in a way that preserves transitivity. The rules for \sqsubseteq_{\leftarrow} are also shown: like $\sqsubseteq_{\rightarrow}$, they allow for reductions before comparing with structural precision, but they also allow backwards steps, fulfilling Lemma 5.9. We only allow backwards-steps for the less precise term, since backwards steps on the left-hand are admissible by Lemma 5.13.

$$\begin{array}{c}
\text{DIAGREFL} \\
\frac{\Gamma_1 | \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\alpha} \mathbf{t}_2 \quad \Gamma_1 | \Gamma_2 \vdash \mathbf{t}'_1 \sqsubseteq_{\alpha} \mathbf{t}'_2}{\Gamma_1 | \Gamma_2 \vdash \mathbf{refl}(\mathbf{t}_1)_{\mathbf{t}'_1} \cong_{\alpha} \mathbf{refl}(\mathbf{t}_2)_{\mathbf{t}'_2}} \\
\text{DIAGCOMP} \\
\frac{\Gamma_1 | \Gamma_2 \vdash \mathbf{T}_1 \sqsubseteq_{\alpha} \mathbf{T}_2 \quad \Gamma_1 | \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\alpha} \mathbf{t}_2 \quad \Gamma_1 | \Gamma_2 \vdash \mathbf{t}'_1 \sqsubseteq_{\alpha} \mathbf{t}'_2}{\Gamma_1 | \Gamma_2 \vdash \mathbf{t}_1 \&_{\mathbf{T}_1} \mathbf{t}'_1 \sqsubseteq_{\alpha} \mathbf{t}_2 \&_{\mathbf{T}_2} \mathbf{t}'_2} \\
\text{PRECCOMPL} \quad \text{PRECCOMPR} \quad \text{CSTCOMP(L,R)} \quad \text{CSTCOMPDIAG} \\
\frac{\Gamma_1 | \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\alpha} \mathbf{t}_3}{\Gamma_1 | \Gamma_2 \vdash \mathbf{t}_1 \&_{\mathbf{T}} \mathbf{t}_2 \sqsubseteq_{\alpha} \mathbf{t}_3} \quad \frac{\Gamma_1 | \Gamma_2 \vdash \mathbf{t}_2 \sqsubseteq_{\alpha} \mathbf{t}_3}{\Gamma_1 | \Gamma_2 \vdash \mathbf{t}_1 \&_{\mathbf{T}} \mathbf{t}_2 \sqsubseteq_{\alpha} \mathbf{t}_3} \quad \frac{\mathbf{t}_1 \cong_{\alpha} \mathbf{t}_3 \quad \mathbf{t}_2 \cong_{\alpha} \mathbf{t}_3}{\mathbf{t}_1 \&_{\mathbf{T}} \mathbf{t}_2 \cong_{\alpha} \mathbf{t}_3} \quad \frac{\mathbf{t}_1 \cong_{\alpha} \mathbf{t}'_1 \quad \mathbf{t}_2 \cong_{\alpha} \mathbf{t}'_2}{\mathbf{t}_1 \&_{\mathbf{T}} \mathbf{t}_2 \cong_{\alpha} \mathbf{t}'_1 \&_{\mathbf{T}'} \mathbf{t}'_2} \\
\text{PRECCOMPSTR} \quad \text{PRECCOMPSTEPL} \quad \text{PRECCOMPSTEPR} \quad \text{PRECCOMPSTEPBACK} \\
\frac{\Gamma_1 | \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\alpha} \mathbf{t}_2}{\Gamma_1 | \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\alpha}^{\leftarrow} \mathbf{t}_2} \quad \frac{\mathbf{t}_1 \longrightarrow^* \mathbf{t}'_1 \quad \Gamma_1 | \Gamma_2 \vdash \mathbf{t}'_1 \sqsubseteq_{\alpha}^{\leftarrow} \mathbf{t}_2}{\Gamma_1 | \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\alpha}^{\leftarrow} \mathbf{t}_2} \quad \frac{\mathbf{t}_2 \longrightarrow^* \mathbf{t}'_2 \quad \Gamma_1 | \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\alpha}^{\leftarrow} \mathbf{t}'_2}{\Gamma_1 | \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\alpha}^{\leftarrow} \mathbf{t}_2} \quad \frac{\mathbf{t}_2 \longrightarrow^* \mathbf{t}_2 \quad \Gamma_1 | \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\alpha}^{\leftarrow} \mathbf{t}'_2}{\Gamma_1 | \Gamma_2 \vdash \mathbf{t}_1 \sqsubseteq_{\alpha}^{\leftarrow} \mathbf{t}_2}
\end{array}$$

Fig. 11. CastEq Precision and Consistency rules for composition

Last are the rules for (static) consistency for composition. Recall that §5.2 requires that reducing precision preserves consistency. Since composition is as precise as both its arguments, $\mathbf{t}_1 \&_{\mathbf{T}} \mathbf{t}_2 \cong_{\alpha} \mathbf{t}_3$ should imply that \mathbf{t}_1 and \mathbf{t}_2 are both consistent with \mathbf{t}_3 . We conjecture that composition is a (semantic) greatest lower bound, which would mean that errors in composing witnesses are never flagged earlier than necessary. For this to hold, the composition of two terms must be consistent with everything that is consistent with both of those two terms. Our composition consistency rules in Fig. 11 express this: CSTCOMPL and CSTCOMPR ensure that $\mathbf{t}_1 \&_{\mathbf{T}} \mathbf{t}_2$ is consistent with exactly the terms that are consistent with both \mathbf{t}_1 and \mathbf{t}_2 .

With consistency fully defined, the difference between static and dynamic consistency is now clear: two terms that share a non-error lower-bound may be statically inconsistent if they differ only in neutral terms. Variables are only statically consistent with themselves (CSTVAR) or ? (CSTUNKL,R). However, for any two variables \mathbf{x} and \mathbf{y} , $\mathbf{x} \&_{\mathbf{T}} \mathbf{y}$ is a non-error term that is as precise as both, as given by PRECCOMP(L,R). This disconnect between precision and consistency is justified by the criteria of §5.2: we show below that reducing precision preserves consistency. The separation of static and dynamic consistency enables the gradual guarantees and conservatively embedding CIC while maintaining static equivalences.

6.3 Fulfilling The Criteria

GEq is now fully defined: we have defined the precision relations \sqsubseteq_{α} and $\sqsubseteq_{\alpha}^{\leftarrow}$, and the consistency relation \cong_{α} , to instantiate \sqsubseteq_{α} , $\sqsubseteq_{\alpha}^{\leftarrow}$ and \cong_{α} . We now establish that these relations fulfill the criteria of §5.2. We give the intuition behind some of the cases that are new compared to GCIC. Full proofs can be found in the appendix of the extended technical report [Eremondi et al. 2022].

- **Immediate Results:** Proving reflexivity of \sqsubseteq_{α} (Lemma 5.4) is a straightforward induction. The rules PRECCOMPL and PRECCOMPR make the composition of terms as precise as either term, proving Lemma 5.7. The closure of $\sqsubseteq_{\alpha}^{\leftarrow}$ under convertibility is built into its definition, proving Lemma 5.9. DIAGCAST gives that casts are monotone, proving Lemma 5.11. The monotonicity of substitution (Lemma 5.12) is proved with a straightforward induction, relying on presynthesis preserving types under substitution. The remaining diagonal rules give that \sqsubseteq_{α} has all structural rules, fulfilling Lemma 5.15.
- **Composition Safety** (Lemma 5.5) For progress, each composition of two canonical forms of the same type has a reduction, where either the heads match and the arguments are composed, or an

error is produced. If one of the composed terms is not a canonical form, then either (1) one of the composed terms can reduce, (2) one term is a $?_T$ or \bar{U}_T where T is not a function or equality type, and we can reduce with REDCOMPUNK or REDCOMPERR , or (3) one of the composed terms is neutral, and hence the composition is neutral. For preservation, either the result is immediate, or casts are inserted to ensure that types are preserved.

- **Composition Confluence** (Lemma 5.6) $\text{REDCOMPUNK}(L,R)$ ensures that composing with $?_T$ only reduces when $?_T$ cannot reduce, avoiding a “diamond” problem.
- **Precision Transitive** (Lemma 5.8): We prove this *after* monotonicity of reduction, which lets us prove that precision-related types have precision-related terms, which is necessary to fulfill premises on terms’ types, such as in CASTL and CASTR . The rest is straightforward induction.
- **Static Consistency** (Lemma 5.10) The GEq rules not present in GCIC are for equality, which are trivially handled, and consistency rules for composition, for which the result vacuously holds since composition is not present in the static language.
- **Monotonicity of Reduction** (Lemma 5.13): The key fact is that, since PRECCOMPL and PRECCOMPR only have composition on the left, all the inversions in the GCIC proofs are still valid for GEq . The interesting case is when precision is derived using PRECCOMPL (PRECCOMPR is symmetric), and the composition reduces. The result is trivial for $\text{REDCOMPUNK}(L,R)$ and $\text{REDCOMPERR}(L,R)$. For the remaining cases, two terms with the same head are being composed, and the result is either \bar{U} or another term with the same head. When \bar{U} is produced the result is trivial. When a term with the same head is produced, the PRECCOMPL can be used with the appropriate diagonal rule. In the case that casts are present in the result of composition, CASTL is used. The other notable case is when \mathbf{J} reduces, where the result is derived using DIAGCAST .
- **Consistency Monotone** (Lemma 5.14) We first show that consistency is monotone on the left, then prove that it is symmetric to obtain monotonicity for both arguments. The case when \cong_α is derived with CSTCOMPR or CSTCASTR must be handled specially, since they each take an operand that can be any term. The trick is to unwrap the chain of CSTCOMPR and CSTCASTR uses, use the induction hypothesis on the contained derivation, then re-apply CSTCOMPR and CSTCASTR in the same order to obtain the result. When precision is derived with PRECCOMPL or PRECCOMPR , then consistency must have either been derived with CSTCOMPDIAG , in which case the result follows from the induction hypothesis, or with CSTCOMPL or (symmetrically) CSTCOMPR . For CSTCOMPL , the premise gives that both composed terms are consistent with the right-hand term, yielding our result. The remaining cases are straightforward.

7 DISCUSSION

7.1 Extensions Enabled by Equality

In addition to catching the kinds of bugs discussed in §2, we show some benefits of having propositional equality in GEq . Three new language features can be encoded using propositional equality, without augmenting the cast calculus: empty types, Axiom K, and indexed inductive type families. For type families, we discuss some limitations of the approach and workarounds for these limitations, showing how our cast calculus is expressive enough to pave the way for future improvements.

7.1.1 The Empty Type. Just as the gradual \mathbf{J} needed computation, eliminating the empty type has computational content in a gradual language. In static languages, the empty type `Empty` has no closed values, so either `Empty` contains no terms, or (for logically inconsistent languages) any such terms are non-terminating. The elimination function $\text{exfalse} : (X : \text{Type}_i) \rightarrow \text{Empty} \rightarrow X$ produces a result of any type, given a value of the empty type. In a gradual language, however, $?$ and \bar{U} can be used at any type, including the empty type. So a gradual `exfalse` must produce a value of type X .

We again follow the goal of dynamically tracking constraints expressed by types. For the empty type, a value of type $f : T \rightarrow \text{Empty}$ encodes the constraint that T should be impossible, and a branch built using `exfalse` should be unreachable. If f is applied to $t : \text{Empty}$, created using `?` or casts, then the constraint has been violated, and an error should be raised.

We can encode this behavior by defining `Empty` to be $\text{true} =_{\mathbb{B}} \text{false}$, and `exfalse` to be $\lambda X. \lambda t. \mathbf{J} (\mathbf{b}. \text{if } \mathbf{b} \mathbb{B} X) \text{ true } t$. The key is `?Empty` and `UEmpty` both evaluate to $\text{refl}(U_{\mathbb{B}})_{\text{true} \cong \text{false}}$. So the only value of type `Empty` is a dynamic type error. Likewise, the eliminator `exfalse` casts t to type U_{Type_e} , before casting it to type X , so the result is always U_X . Without adding any features to `CastEq`, the bug-finding described in §2 handles constraints encoded as logical negation.

7.1.2 Axiom K. Because $?_{t_1 =_{\mathbb{T}} t_2}$ steps to $\text{refl}(t_1 \ \&_{\mathbb{T}} \ g_2)_{t_1 \cong t_2}$, `GEq` is in the class of dependently typed languages where `refl` is the only constructor for equality. Composition can be used to derive a (gradual) proof of this uniqueness, even though no such proof can be derived in most static type theories [[Hofmann and Streicher 1998](#)]:

$$\begin{aligned} K : (x : T) \rightarrow (\text{pf} : x =_{\mathbb{T}} x) \rightarrow \text{pf} =_{x =_{\mathbb{T}} x} \text{refl}(x)_{x \cong x} \\ K \ x \ \text{pf} = \text{refl}(\text{pf} \ \&_{x =_{\mathbb{T}} x} \ \text{refl}(x)_{x \cong x})_{\text{pf} \cong \text{refl}(x)_{x \cong x}} \end{aligned}$$

Axiom K can be used to prove that all equality proofs of a given type are equal [[Streicher 1993](#)], so our proof-irrelevant `J` principle does not lose any expressivity, since any types parameterized by an equality proof can be rewritten with `K`. Also, Axiom K allows for conventional dependent pattern matching to be elaborated into inductive eliminators [[Goguen et al. 2006](#)], providing a lightweight alternative to the cumbersome `ind` form. The combination of Axiom K and function extensionality suggests a connection to Observational Type Theory [[Altenkirch et al. 2007](#); [Pujet and Tabareau 2022](#)] that warrants future exploration.

7.1.3 Inductive Types. [McBride \[2000\]](#) describes how, using propositional equality, indexed inductive families can be encoded. The main idea is, instead of having each constructor return different indices, each index is a parameter, and each constructor takes an equality proof that the parameter has the desired value. In the elimination principle, `J` is used to rewrite the type of the returned value using the stored equality. Consider how the classic vector type is transformed:

$$\begin{aligned} \text{data } \text{Vec} (X : \text{Type}) : (n : \mathbb{N}) \rightarrow \text{Type} \text{ where} & \quad \text{data } \text{Vec}' (X : \text{Type}) (n : \mathbb{N}) : \text{Type} \text{ where} \\ \text{Nil} : \text{Vec } X \ 0 & \quad \text{Nil}' : (n =_{\mathbb{N}} 0) \rightarrow \text{Vec } X \ n \\ \text{Cons} : X \rightarrow \text{Vec } X \ n \rightarrow \text{Vec } X \ (1+n) & \quad \text{Cons}' : (z : \mathbb{N}) \rightarrow X \rightarrow \text{Vec } X \ z \\ & \quad \rightarrow n =_{\mathbb{N}} (1+z) \rightarrow \text{Vec } X \ n \end{aligned}$$

This transformation gives a low-overhead way to incorporate indexed inductive families with gradual dependent types. Since no extensions to `CastEq` are required, the safety and gradual guarantee results from §5 apply. The constructors take equality proofs, so violations of those equalities raise dynamic type errors.

However, the approach is limited in its ability to eagerly detect errors. The problem is that dynamic consistency is fundamentally not transitive, since otherwise all types are consistent through `?`. Members of inductive types are essentially trees, and equality constraints track constraints at each level of the tree, but consistency across the entire tree is not ensured. The witnesses track the evolution of type information across time, but not across space. Consider $\text{Cons}' \ ?_{\mathbb{N}} \ \text{true} \ (\text{Nil}' \ ?_{\mathbb{N}} \ \text{refl}(0)_{0 =_{\mathbb{N}} 0}) \ \text{refl}(2)_{2 =_{\mathbb{N}} 2} : \text{Vec}' \ \mathbb{B} \ 2$, a vector with one element, whose type says it has length 2. Constructing this vector raises no run-time type errors. At each level, the equality proof is correct: `0` is consistent with `?N`, and `2` is consistent with `1 + ?N`. Gradually, the non-transitivity means that imprecision at each level can cause disconnects between levels.

Thankfully, CastEq is expressive enough to encode a solution to this problem. By having composition as an operator in the language, one can define so-called “smart constructors” that have the same types as the normal constructors, but that access the equality proofs stored in the previous level of the tree when constructing new ones. For example, using **J** we can write $\text{cong1} : (\mathbf{m} \ \mathbf{n} : \mathbb{N}) \rightarrow \mathbf{m} =_{\mathbb{N}} \mathbf{n} \rightarrow 1 + \mathbf{m} = 1 + \mathbf{n}$, which can be used in a “smart” **Cons**:

$$\text{smartCons } \mathbf{z} \ \mathbf{h} \ \mathbf{t} \ \mathbf{eq} = \text{Cons } \mathbf{z} \ \mathbf{h} \ \mathbf{t} \ (\mathbf{eq} \ \&_{(1+\mathbf{z})=_{\mathbb{N}} \mathbf{n}} \ ((1+\mathbf{z}) =_{\mathbb{N}} \mathbf{n} \Leftarrow 1+\mathbf{z} =_{\mathbb{N}} 1+\mathbf{z})(\text{cong1} \ (\text{wit} \ \mathbf{t}))))$$

where $\text{wit} : \text{Vec } X \ \mathbf{z} \rightarrow \mathbf{z} =_{\mathbb{N}} \mathbf{z}$

$$\text{wit} \ (\text{Nil} \ \mathbf{eq}) = \langle \mathbf{z} =_{\mathbb{N}} \mathbf{z} \Leftarrow 0 =_{\mathbb{N}} \mathbf{z} \rangle \mathbf{eq} \quad | \quad \text{wit} \ (\text{Cons } \mathbf{x} \ \mathbf{h} \ \mathbf{t} \ \mathbf{eq}) = \langle \mathbf{z} =_{\mathbb{N}} \mathbf{z} \Leftarrow 1+\mathbf{x} =_{\mathbb{N}} \mathbf{z} \rangle \mathbf{eq}$$

When **smartCons** is used in place of **Cons**’, the witness $\text{refl}(0)_{\cdot, ?_{\mathbb{N}} \cong 0}$ is transformed to $\text{refl}(1)_{\cdot, ?_{\mathbb{N}} \cong ?_{\mathbb{N}}}$, which produces an error when cast to $1+?_{\mathbb{N}} =_{\mathbb{N}} 2$, since $1 \ \&_{\mathbb{N}} \ 2 \rightsquigarrow \mathcal{U}_{\mathbb{N}}$. Formalizing the general version of this approach is beyond the scope of this paper, but it shows how having composition as an operator enables more detailed manipulation of run-time type information.

7.2 Future Work

7.2.1 Termination and Approximate Normalization. As presented, GEq has undecidable type checking. Precision and CastEq typing are not obviously decidable. Thankfully, we do not need to decide precision typing to decide GEq typing, since initial witnesses are always valid by reflexivity. Deciding GEq typing is straightforward, except for normalization: some terms do not terminate and consistency compares modulo reduction. In GCIC, [Lennon-Bertrand et al. \[2022\]](#) show that termination can be obtained by sacrificing the gradual guarantees, or by restricting universes so that they are not closed under function types. While useful for type theory, these sacrifices remove reasoning principles or reduce expressivity, respectively, that make them unsuitable for programming.

[Eremondi et al. \[2019\]](#) propose *approximate normalization*, with different semantics for compile-time normalization of types and run-time evaluation of terms. At compile-time, when missing type information means that termination cannot be guaranteed, $?$ is produced as an approximation. Run-time evaluation uses no approximations, so expressivity is not lost. We conjecture that approximate normalization can be easily added to GEq. The challenge is finding a suitable termination argument, since [Eremondi et al. \[2019\]](#) provide a proof that does not apply to inductive types. The syntactic-model strategy of [Lennon-Bertrand et al. \[2022\]](#) can likely be adapted. Also, a decision must be made about whether approximate or exact normalization should be used for run-time witness calculations: since they are opaque to the programmer, witnesses may obscure the causes of non-termination.

7.2.2 Conjectures: EP Pairs, Composition, Blame and Full Abstraction. [Lennon-Bertrand et al. \[2022\]](#) prove a stronger property than the gradual guarantees for GCIC. They show that casts between precision-related types form an embedding-projection (EP) pair [[New and Ahmed 2018](#)], so that increasing then decreasing precision produces the same result modulo errors, and decreasing then increasing precision produces an observationally-equivalent result. While the gradual guarantees are helpful, they are satisfied by trivial languages where every cast produces $?$. Showing the EP pair property would prove that casts in GEq never lose run-time information, giving more confidence in its ability to dynamically track constraints. We conjecture that GEq fulfills the EP pair property, but suspect novel proof techniques are needed to handle witness proofs. We also conjecture that composition computes the greatest lower-bound for semantic precision, so that each type forms a true semi-lattice. This would establish that witness composition never prematurely raises dynamic errors, since two witnesses would compose to \mathcal{U} only when all other options are impossible.

Another desirable property of gradual languages is a blame theorem [Wadler and Findler 2009], stating that, for a dynamic cast error, the less-precise type is always blamed. GEq has no notion of blame, but we conjecture that the techniques of Zalewski et al. [2020] could be adapted to GEq.

Finally, we conjecture that there is a variant of CIC whose embedding into GEq is fully abstract, meeting the criteria Jacobs et al. [2021] set out for gradual languages. Intuitively, we can form equalities between extensionally-equal functions, and use those to cast between types indexed by them, so any property of a function should apply to one that is extensionally-equal. Full abstraction guarantees that all static equivalences hold in GEq, giving the programmer more tools with which to reason about their code. Proving full abstraction for non-dependently typed gradual languages is a recent development, so more investigation is needed to adapt these techniques to dependent types. The usual technique for full abstraction is to simulate the target language in the source, so every target context can be translated into a source context that is unable to distinguish the terms. As a consequence, the embedded variant of CIC must have capabilities for non-termination added.

7.3 Related Work

Gradual Approaches to Equality: GRIP [Maillard et al. 2022] extends CastCIC with propositional equality, but unlike GEq, this propositional equality is in a separate sort, the types of which have all members definitionally equal. This propositional layer has no imprecision, and contains features for catching dynamic type errors in the gradual layer. GRIP features an *internal* notion of precision: while not all terms obey the gradual guarantees, self-precise terms do, and guarantees for such terms are available as a theorem in the propositional layer. GEq provides reasoning about equality *within* gradual programs, whereas GRIP provides a layer for reasoning *about* gradual programs. GRIP is based on observational type theory (OTT), and hence supports fully static proofs of function extensionality and Axiom K. GEq contains only gradual proofs of these. We conjecture that, like in GRIP, OTT could replace BCIC as the underlying static language for GEq.

Our work also relates to that of Lemay [2022], which presents a dependently typed language where all definitional equality checks are deferred to run time. Like in GCIC and GEq, casts are inserted when a term's synthesized type differs from the type against which it is checked. However, there is no unknown term/type, and no requirement that synthesized and checked-against types be consistent, so errors are raised only when safety is directly violated. Like in GEq, function equality is checked extensionally: casts accumulate on unapplied functions, and arguments are cast when cast-containing functions are applied, similar to composing functions in GEq. The work emphasizes clear error messages, and unlike GEq, features a rich notion of blame.

Flexible Dependent Types: GEq builds on a long line of work mixing dynamic and static enforcement of specifications. Ou et al. [2004] support mixed static and dynamic checking of boolean-valued properties, and Lehmann and Tanter [2017] provide gradual typing for refinement types. Similarly, Tanter and Tabareau [2015] develop a system of casts for Coq, using an unsound axiom to represent type errors. These casts worked for values of subset types, i.e. a value paired with a proof that some boolean-valued function returns true for that value, but not general inductive types. Osera et al. [2012] present *dependent interoperability* for principled mixing of dependently typed and non-dependently typed programs. Dependent interoperability was extended by Dagand et al. [2016, 2018], who provide a general mechanism for lifting higher-order programs to the dependently typed setting. All of these approaches presuppose separate simple and dependent versions of types, related by boolean-valued predicates. Our composition of witnesses provides similar checks, but by keeping witnesses, types need not be reformulated in terms of subset types or boolean predicates.

Acknowledgments: We thank Felipe Bañados Schwerter, Paulette Koronkevich, Jonathan Chan, Joanne Traves, and the anonymous reviewers for their feedback on this work.

REFERENCES

- Thorsten Altenkirch, Conor McBride, and Wouter Swierstra. 2007. Observational Equality, Now!. In Proceedings of the 2007 Workshop on Programming Languages Meets Program Verification (Freiburg, Germany) (PLPV '07). ACM, New York, NY, USA, 57–68. <https://doi.org/10.1145/1292597.1292608>
- Felipe Bañados Schwerter, Alison M. Clark, Khurram A. Jafery, and Ronald Garcia. 2021. Abstracting Gradual Typing Moving Forward: Precise and Space-Efficient. *Proc. ACM Program. Lang.* 5, POPL, Article 61 (Jan. 2021), 28 pages. <https://doi.org/10.1145/3434342>
- Edwin Brady. 2017. Type-driven development with Idris. Manning. <https://www.manning.com/books/type-driven-development-with-idris>
- Pierre-Évariste Dagand, Nicolas Tabareau, and Éric Tanter. 2016. Partial Type Equivalences for Verified Dependent Interoperability. In Proceedings of the 21st ACM SIGPLAN Conference on Functional Programming (ICFP 2016). ACM Press, Nara, Japan, 298–310. <https://doi.org/10.1145/2951913.2951933>
- Pierre-Évariste Dagand, Nicolas Tabareau, and Éric Tanter. 2018. Foundations of dependent interoperability. *Journal of Functional Programming* 28 (2018), e9. <https://doi.org/10.1017/S0956796818000011>
- Joseph Eremondi, Ronald Garcia, and Éric Tanter. 2022. Propositional Equality for Gradual Dependently Typed Programming (Extended Technical Report). <https://doi.org/10.48550/ARXIV.2205.01241>
- Joseph Eremondi, Éric Tanter, and Ronald Garcia. 2019. Approximate Normalization for Gradual Dependent Types. *Proc. ACM Program. Lang.* 3, ICFP, Article 88 (July 2019), 30 pages. <https://doi.org/10.1145/3341692>
- Robert Bruce Findler and Matthias Felleisen. 2002. Contracts for Higher-order Functions. In Proceedings of the Seventh ACM SIGPLAN International Conference on Functional Programming (Pittsburgh, PA, USA) (ICFP '02). ACM, New York, NY, USA, 48–59. <https://doi.org/10.1145/581478.581484>
- Ronald Garcia, Alison M. Clark, and Éric Tanter. 2016. Abstracting Gradual Typing. In Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (St. Petersburg, FL, USA) (POPL '16). ACM, New York, NY, USA, 429–442. <https://doi.org/10.1145/2837614.2837670>
- Healfdene Goguen, Conor McBride, and James McKinna. 2006. Eliminating Dependent Pattern Matching. Springer Berlin Heidelberg, Berlin, Heidelberg, 521–540. https://doi.org/10.1007/11780274_27
- Martin Hofmann and Thomas Streicher. 1998. The groupoid interpretation of type theory. Twenty-five years of constructive type theory (Venice, 1995) 36 (1998), 83–111. <https://doi.org/10.1093/oso/9780198501275.003.0008>
- Koen Jacobs, Amin Timany, and Dominique Devriese. 2021. Fully Abstract from Static to Gradual. *Proc. ACM Program. Lang.* 5, POPL, Article 7 (Jan. 2021), 30 pages. <https://doi.org/10.1145/3434288>
- Nico Lehmann and Éric Tanter. 2017. Gradual Refinement Types. In Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (Paris, France) (POPL 2017). ACM, New York, NY, USA, 775–788. <https://doi.org/10.1145/3009837.3009856>
- Mark Lemay. 2022. A Dependently Typed Programming Language With Dynamic Equality. Ph. D. Dissertation. Boston University. <https://github.com/marklemay/thesis>
- Meven Lennon-Bertrand. 2021. Complete Bidirectional Typing for the Calculus of Inductive Constructions. In 12th International Conference on Interactive Theorem Proving (ITP 2021) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 193), Liron Cohen and Cezary Kaliszzyk (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik. <https://doi.org/10.4230/LIPIcs.ITP.2021.24>
- Meven Lennon-Bertrand, Kenji Maillard, Nicolas Tabareau, and Éric Tanter. 2022. Gradualizing the Calculus of Inductive Constructions. *ACM Trans. Program. Lang. Syst.* 44, 2, Article 7 (apr 2022), 82 pages. <https://doi.org/10.1145/3495528>
- Kenji Maillard, Meven Lennon-Bertrand, Nicolas Tabareau, and Éric Tanter. 2022. A Reasonably Gradual Type Theory. *Proc. ACM Program. Lang.* 6, ICFP (2022). <https://doi.org/10.1145/3547655> Preprint: <https://hal.inria.fr/hal-03596652/>.
- Per Martin-Löf. 1975. About Models for Intuitionistic Type Theories and the Notion of Definitional Equality. In Proceedings of the Third Scandinavian Logic Symposium, Stig Kanger (Ed.). Studies in Logic and the Foundations of Mathematics, Vol. 82. Elsevier, 81 – 109. [https://doi.org/10.1016/S0049-237X\(08\)70727-4](https://doi.org/10.1016/S0049-237X(08)70727-4)
- Per Martin-Löf. 1982. Constructive Mathematics and Computer Programming. In Logic, Methodology and Philosophy of Science VI, L. Jonathan Cohen, Jerzy Łoś, Helmut Pfeiffer, and Klaus-Peter Podewski (Eds.). Studies in Logic and the Foundations of Mathematics, Vol. 104. Elsevier, 153–175. [https://doi.org/10.1016/S0049-237X\(09\)70189-2](https://doi.org/10.1016/S0049-237X(09)70189-2)
- Conor McBride. 2000. Dependently typed functional programs and their proofs. Ph. D. Dissertation. University of Edinburgh, UK. <http://hdl.handle.net/1842/374>
- Conor McBride. 2002. Elimination with a Motive. In Types for Proofs and Programs (Berlin, Heidelberg, 2002) (Lecture Notes in Computer Science), Paul Callaghan, Zhaohui Luo, James McKinna, Robert Pollack, and Robert Pollack (Eds.). Springer, 197–216. https://doi.org/10.1007/3-540-45842-5_13
- Max S. New and Amal Ahmed. 2018. Graduality from Embedding-Projection Pairs. *Proc. ACM Program. Lang.* 2, ICFP, Article 73 (July 2018), 30 pages. <https://doi.org/10.1145/3236768>

- Peter-Michael Osera, Vilhelm Sjöberg, and Steve Zdancewic. 2012. Dependent Interoperability. In Proceedings of the Sixth Workshop on Programming Languages Meets Program Verification (Philadelphia, Pennsylvania, USA) (PLPV '12). ACM, New York, NY, USA, 3–14. <https://doi.org/10.1145/2103776.2103779>
- Xinming Ou, Gang Tan, Yitzhak Mandelbaum, and David Walker. 2004. Dynamic Typing with Dependent Types. In Exploring New Frontiers of Theoretical Informatics, Jean-Jacques Levy, Ernst W. Mayr, and John C. Mitchell (Eds.). Springer US, Boston, MA, 437–450. https://doi.org/10.1007/1-4020-8141-3_34
- Loïc Pujet and Nicolas Tabareau. 2022. Observational Equality: Now for Good. Proc. ACM Program. Lang. 6, POPL, Article 32 (jan 2022), 27 pages. <https://doi.org/10.1145/3498693>
- Jeremy Siek, Peter Thiemann, and Philip Wadler. 2015a. Blame and Coercion: Together Again for the First Time. In Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation (Portland, OR, USA) (PLDI '15). Association for Computing Machinery, New York, NY, USA, 425–435. <https://doi.org/10.1145/2737924.2737968>
- Jeremy G. Siek and Tianyu Chen. 2021. Parameterized cast calculi and reusable meta-theory for gradually typed lambda calculi. Journal of Functional Programming 31 (2021), e30. <https://doi.org/10.1017/S0956796821000241>
- Jeremy G. Siek and Walid Taha. 2006. Gradual Typing for Functional Languages. In Scheme and Functional Programming Workshop. 81–92. <http://scheme2006.cs.uchicago.edu/scheme2006.pdf>
- Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, and John Tang Boyland. 2015b. Refined Criteria for Gradual Typing. In 1st Summit on Advances in Programming Languages (SNAPL 2015) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 32), Thomas Ball, Rastislav Bodik, Shriram Krishnamurthi, Benjamin S. Lerner, and Greg Morrisett (Eds.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 274–293. <https://doi.org/10.4230/LIPIcs.SNAPL.2015.274>
- Jeremy G. Siek and Philip Wadler. 2010. Threesomes, with and Without Blame. In Proceedings of the 37th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (Madrid, Spain) (POPL '10). ACM, New York, NY, USA, 365–376. <https://doi.org/10.1145/1706299.1706342>
- Thomas Streicher. 1993. Investigations into intensional type theory. Ph. D. Dissertation. Ludwig Maximilian Universität. <https://www2.mathematik.tu-darmstadt.de/~streicher/HabilStreicher.pdf> Habilitation thesis.
- M. Takahashi. 1995. Parallel Reductions in λ -Calculus. Information and Computation 118, 1 (1995), 120–127. <https://doi.org/10.1006/inco.1995.1057>
- Éric Tanter and Nicolas Tabareau. 2015. Gradual Certified Programming in Coq. In Proceedings of the 11th Symposium on Dynamic Languages (Pittsburgh, PA, USA) (DLS 2015). ACM, New York, NY, USA, 26–40. <https://doi.org/10.1145/2816707.2816710>
- Philip Wadler and Robert Bruce Findler. 2009. Well-Typed Programs Can't Be Blamed. In Proceedings of the 18th European Symposium on Programming Languages and Systems: Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009 (York, UK) (ESOP '09). Springer-Verlag, Berlin, Heidelberg, 1–16. https://doi.org/10.1007/978-3-642-00590-9_1
- A.K. Wright and M. Felleisen. 1994. A Syntactic Approach to Type Soundness. Information and Computation 115, 1 (1994), 38–94. <https://doi.org/10.1006/inco.1994.1093>
- Jakub Zalewski, James McKinna, J. Garrett Morris, and Philip Wadler. 2020. λ dB: Blame tracking at higher fidelity. <https://wgt20.irif.fr/wgt20-final98-acmpaginated.pdf> First ACM SIGPLAN Workshop on Gradual Typing 2020, WGT 2020 ; Conference date: 19-01-2020 Through 25-01-2020.