Dissensus AI Working Paper Series
DAI-2603

Identity is Irreducibly Relational

A Critique of Primitive Identity from ZFC to Homotopy Type Theory

Murad Farzulla[1,2,*]

[1]Dissensus AI, London, UK [2]King's College London, London, UK

[*]Correspondence: murad@dissensus.ai ORCID: 0009-0002-7164-8704

January 2026

Abstract

This paper argues that identity is irreducibly relational: the statement $A = A$ presupposes that A is defined, and definition requires distinction from a background. I develop this thesis at three levels: *conceptual* (definition requires distinction), *formal* (examining how set-theoretic and type-theoretic foundations treat identity), and *historical* (engaging the literature from Leibniz through Kripke). Against the standard view that identity is primitive in first-order logic, I argue that this primitiveness reflects a genuine conceptual difficulty that modern foundations—particularly Homotopy Type Theory and Univalent Foundations—have begun to resolve by treating identity as constituted by structural equivalence. The extensionality axiom of set theory already makes identity relational for sets; the univalence axiom generalizes this insight. I conclude that the trajectory of foundational mathematics vindicates a relational conception of identity, with implications for metaphysics, philosophy of mathematics, and the "hard problems" that arise when transformation fails to preserve structure.

Keywords: identity, reference, relational ontology, ZFC, Homotopy Type Theory, univalence, structuralism

1 Introduction: The Presupposition in $A = A$

The law of identity—$A = A$—appears to be the most primitive and unassailable claim in logic. It seems to require no justification, admit no exception, and presuppose nothing.

This paper argues otherwise. The statement $A = A$ presupposes that A is defined, and definition is not primitive. To define A is to distinguish it from what it is not—which requires reference to something external to A. Identity therefore presupposes referential structure, and the law of identity is derivative rather than foundational.

This claim must be carefully distinguished from adjacent positions. It is not Geach's (1967) relative identity, which denies absolute identity in favour of sortal-relative identity. I maintain that identity is absolute but *relationally constituted*: two things that are the same F are simply the same, but what makes them so is their relational position, not an intrinsic property of "sameness." Nor is it mere sortalism in the sense of Wiggins (2001), who treats sortals as governing the *application* of identity criteria without relativizing identity itself. I argue that the ineliminable role of sortals reveals something constitutive: if identity cannot even be applied without relational apparatus, identity is not independent of that apparatus. Finally, the thesis is not epistemological. The claim is not that we need reference to *know* what A is, but that A requires reference to *be defined at all*. This is an ontological thesis about what identity consists in.

The argument proceeds as follows. Section 2 formalizes the thesis: definition requires distinction, and distinction is relational. Section 3 situates this claim within the philosophical literature on identity, from Leibniz through Kripke to Fine's grounding framework. Sections 4–6 trace the treatment of identity across three foundational programmes—set theory, category theory, and Homotopy Type Theory—showing a progressive shift from primitive to relationally constituted identity. Section 7 demonstrates that HoTT dissolves the identity problem that has plagued mathematical structuralism. Section 8 surveys convergent developments in analytic metaphysics. Section 9 addresses objections, and Sections 10–11 draw out consequences and synthesize the convergence argument.

2 The Formal Thesis

Definition 2.1 (Referential Set)**.** Let $\delta(A,x)$ be a primitive relation of distinguishability between A and x. For any entity A, the **referential set** $R(A)$ is the set of entities from which A is distinguishable:

$$R(A) = \{x : \delta(A,x)\}$$

Axiom 2.2 (Irreflexivity of Distinguishability)**.** No entity is distinguishable from itself:

$$\forall A\, \neg\delta(A,A)$$

Axiom 2.3 (Non-Triviality)**.** For any defined entity A, there exists some x from which A is distinguishable:

$$\mathrm{Def}(A) \implies \exists x\, \delta(A,x)$$

Remark 2.4 (Avoiding Circularity). This formulation avoids presupposing identity in the definition of $R(A)$. The earlier version included "$x \neq A$" in the set-builder, which circularly invoked non-identity. Here, the condition $x \neq A$ is *derived* from Axiom 2.2: since $\neg\delta(A,A)$, $A \notin R(A)$ follows automatically. The primitive distinguishability relation δ thus does the work without presupposing the identity relation it aims to ground.

Remark 2.5 (Status of the Framework). The framework presented here is not itself a formal theory within a fixed logical system but a *schema* that can be instantiated within various foundational settings. In ZFC or NBG, δ might be interpreted as a definable relation on the domain of discourse; in HoTT, as a proposition-valued relation on types; in a topos, via the internal logic using a subobject classifier. The schema captures a structural insight—that definition requires distinction—while remaining neutral on which foundation implements it. The "theorems" that follow are therefore schematic: they hold in any instantiation where the axioms are satisfied. This meta-level framing is deliberate. The claim that identity presupposes definition cannot be stated *within* a system that takes identity as primitive—it is a claim *about* such systems, visible only from a metatheoretic vantage point.

Axiom 2.6 (Definition Requires Distinction)**.** An entity A is defined if and only if $R(A) \neq \emptyset$:

$$\mathrm{Def}(A) \iff R(A) \neq \emptyset$$

Together with Axiom 2.3, this entails that defined entities are distinguishable from at least one thing.

This axiom captures a fundamental insight: to define something is to say what it is by saying what it is not. An entity that cannot be contrasted with anything—that is indistinguishable from every other entity and from the background itself—is not defined. It has no referential foothold; the symbol purporting to name it is an empty gesture.

Theorem 2.7 (The Relational Identity Thesis)**.** *The identity claim $A = A$ presupposes that A is defined. Therefore:*

$$(A = A) \implies R(A) \neq \emptyset$$

Identity is not primitive but derivative of referential structure.

Proof. This is a metatheoretic argument, not a derivation within a fixed object-level formal system. The premise "$A = A$ is meaningful" is semantic: it concerns the conditions under which the formula has interpretive content rather than being an uninterpreted string.

Suppose $A = A$ is a meaningful claim. For this claim to have content—for it to say something about A—the symbol A must refer to something in the domain of discourse. That is, A must be defined. By Axiom 2.6, $\mathrm{Def}(A) \iff R(A) \neq \emptyset$. Since A is defined (otherwise $A = A$ would be a meaningless string of symbols, not a proposition), we have $R(A) \neq \emptyset$. Therefore $(A = A) \implies R(A) \neq \emptyset$.

The metatheoretic character of this argument is a feature, not a defect. The thesis concerns what is *presupposed* when identity claims are meaningful—a question that necessarily stands outside any object-level system in which identity is axiomatized. ■

Corollary 2.8 (No Self-Sufficient Identity)**.** *No entity can be self-identical in a meaningful sense without reference to something external to itself. "Pure self-identity"—identity without relation—is not false but contentless.*

3 Identity in the Philosophical Literature

The δ-framework formalizes a claim about the presuppositional structure of identity. This section examines how the philosophical tradition has treated identity—from the standard logical treatment through Leibniz, Kripke, and Fine—and argues that even those who take identity as primitive implicitly rely on relational apparatus.

3.1 The Standard View: Identity as Primitive

In first-order logic with equality, identity is typically taken as primitive. The axioms are:

$$\forall x(x=x) \qquad \text{(Reflexivity)} \tag{1}$$

$$\forall x \forall y(x=y \rightarrow (\phi(x) \rightarrow \phi(y))) \qquad \text{(Substitutivity)} \tag{2}$$

From these, symmetry and transitivity are derivable. But identity itself is "often considered a primitive notion, meaning it is not formally defined, but rather informally said to be 'a relation each thing bears to itself and nothing else' " (Gallois, 2016).

The characterization is **notably circular**: "nothing else" already invokes non-identity. This circularity is typically dismissed as unavoidable for primitive notions. I argue it is not unavoidable but *diagnostic*—it reveals that identity is not genuinely primitive but presupposes a background of distinction.

3.2 Leibniz's Principles

Leibniz proposed two principles:

- **Indiscernibility of Identicals (LL1):** If $x=y$, then for any property F, $Fx \iff Fy$.
- **Identity of Indiscernibles (LL2):** If for any property F, $Fx \iff Fy$, then $x=y$.

LL1 is uncontroversial. LL2 is disputed: Black (1952) offers the thought experiment of two qualitatively indiscernible iron spheres in an otherwise empty universe. If sound, this shows qualitative indiscernibility does not entail numerical identity.

For my purposes, the Leibnizian framework is significant because it *attempts to define identity in terms of property-sharing*—a relational characterization. The debate over LL2 concerns whether this characterization succeeds, not whether identity should be analyzed relationally at all.

3.3 The Hilbert-Bernays Definition

In higher-order logic, identity can be *defined* rather than taken as primitive (Hilbert and Bernays, 1934; Quine, 1986). The question of identity criteria for abstract objects has a longer pedigree, originating in Frege (1884), whose *Grundlagen* remains the locus classicus. Frege's strategy—defining number identity via equinumerosity (Hume's Principle)—already treats identity as constituted by a relational condition rather than grasped directly. His Julius Caesar problem (how to determine whether a number is identical to an arbitrary object) reveals that identity criteria cannot function without a prior specification of domain: the relational apparatus must delimit what counts as a candidate for identification before identity can be assessed. The Hilbert-Bernays definition (anticipated by Leibniz) states:

$$x=y \equiv_{\text{def}} \forall P(Px \leftrightarrow Py)$$

Two objects are identical if and only if they share all properties. This eliminates identity as a primitive in favor of second-order quantification over properties.

One might think this definition already provides a relational account of identity, rendering the δ-framework superfluous. But there is a crucial difference in explanatory level. The Hilbert-Bernays definition tells us *when* two objects are identical—namely, when they are property-indiscernible. It presupposes that x and y are already well-defined terms over which properties can be predicated.

The δ-framework operates at a prior level: it concerns the conditions under which terms *enter the domain of discourse* at all. Before we can ask whether x and y share all properties, x and y must be defined—distinguishable from the background. The question "what grounds the availability of x as a subject of predication?" is not answered by the Hilbert-Bernays definition; it is presupposed.

Thus the two approaches are complementary. Hilbert-Bernays provides a relational criterion for identity among admitted objects. The δ-framework explicates what is required for objects to be admitted in the first place. The relational character of identity operates at both levels.

3.4 Free Logic and Existence Presuppositions

The δ-framework's central insight—that identity claims presuppose that their terms denote—has formal precedent in *free logic*, developed by Lambert and others (Lambert, 1991; Bencivenga, 2002). Free logics modify classical first-order logic to handle non-denoting terms systematically.

In classical logic, the schema $t = t$ (self-identity) is valid for any term t. But this validity presupposes that t denotes—that there exists something to which t refers. Free logics make this presupposition explicit. In *negative free logic*, $t = t$ fails when t is an empty term: "Pegasus = Pegasus" is not true but false or truth-valueless, because there is no Pegasus to be self-identical with.

As Nolt (2020) explains, free logics distinguish between the *outer domain* (all terms the language can form) and the *inner domain* (existing objects). Identity claims are true only when their terms denote members of the inner domain. The classical inference from $t = t$ to $\exists x(x = t)$—which would absurdly prove that anything we can name exists—is blocked.

This is precisely the formal apparatus the δ-framework invokes informally. When I argue that "$A = A$ presupposes that A is defined," I am making a claim that free logic has already formalized: the reflexivity schema requires an existence presupposition. The δ-framework adds the further claim that *being defined* (entering the inner domain) itself requires distinguishability from something—a relational condition.

The connection to free logic strengthens the thesis in two ways. First, it shows that the existence presupposition in identity claims is not a philosophical novelty but a recognized formal phenomenon (Lehmann, 2002). Second, it clarifies the paper's contribution: free logic formalizes *that* identity presupposes existence, while the δ-framework analyzes *what* existence (qua definedness) consists in—namely, relational distinguishability. The frameworks are complementary layers of the same insight.

3.5 Quine: No Entity Without Identity

Quine (1981) famously held that ontological admissibility requires clear individuation criteria: "No entity without identity." But Quine treats identity itself as primitive within first-order logic. His slogan demands criteria for *when* identity holds without analyzing *what* identity is.

This leaves a gap. Quine asks: under what conditions is x identical to y? But this presupposes we already understand what identity means. The question "when are x and y identical?" presupposes an answer to "what would it be for x and y to be identical?"

Quine's own later work suggests he sensed this gap. In "Ontological Relativity," Quine (1968) argues that reference is indeterminate—the question "does 'gavagai' refer to a rabbit, a rabbit-stage, or an undetached rabbit-part?" has no fact of the matter outside a background language. If reference is framework-relative, then identity is too: whether a and b refer to the same thing depends on how the reference scheme individuates objects. This is a stronger concession to the relational thesis than Quine's earlier slogan suggests, since it entails that identity cannot be settled independently of the relational apparatus of a language and its ontology.

3.6 Kripke: Necessary Identity

Kripke (1980) argues that if 'a' and 'b' are rigid designators, then "$a = b$," if true, is necessarily true. This yields the *necessary a posteriori*: "Hesperus = Phosphorus" is necessary (there is no world where

Hesperus is not Phosphorus) but knowable only empirically.

For my thesis, Kripke's examples are revealing. Every case of informative identity (Hesperus/Phosphorus, water/H_2O) involves discovering that two *modes of presentation* converge on the same referent. We never grasp identity directly; we discover it through the convergence of descriptions. This suggests that identity claims are always *mediated* by representational frameworks—consistent with a relational conception.

3.7 Fine: Essence, Ground, and Identity Criteria

Fine (1994) argues that essence cannot be analyzed modally. Socrates is necessarily distinct from the Eiffel Tower, but this distinctness is not *essential* to Socrates—"nothing in Socrates' nature connects him in any special way to the Eiffel Tower." What something *is*—its essence—is prior to what is necessarily true of it.

More directly relevant to our thesis, Fine (2016) argues that identity criteria are *statements of ground*. To give identity conditions for Fs is not merely to state when two Fs are the same, but to say what *grounds* their identity—what makes it the case that they are identical. This framework has profound implications: if identity requires grounding, then identity is not primitive but *derivative*. The identity of Fs is grounded in something more fundamental than identity itself.

Fine's grounding framework supports the relational thesis directly. Consider: what could ground identity if not relational/structural features? The candidates are limited. Either identity is grounded in intrinsic properties (which, as I argue, are themselves relational), or in structural position, or in equivalence relations. All three options make identity derivative of something relational. The alternative—that identity is primitive and ungrounded—conflicts with Fine's insight that identity criteria are genuinely explanatory, not merely extensionally correct.

This connection between essence, ground, and identity has been developed further by Koslicki (2020) and Correia and Skiles (2019), who explore how identity, essence, and grounding interrelate. Carrara and Florio (2020) pursue this line further, arguing that identity criteria serve as epistemic paths to conceptual grounding: they do not merely tell us when Fs are the same but provide the conceptual resources through which we grasp what Fs are. On their account, identity criteria are constitutive of the concepts they govern, which supports the view that identity is not a standalone primitive but depends on the relational apparatus of individuation.

The abstractionist tradition reaches a convergent conclusion from a different direction. Linnebo (2018) argues that abstract objects are "thin"—their existence is a trivial consequence of providing appropriate identity criteria. To be an object, on this view, is to be a possible referent of a singular term, and singular reference is achieved precisely by specifying a criterion of identity. If identity criteria are what bring objects into the ontological fold, then identity is prior to objecthood rather than posterior to it—and since identity criteria are always relational (comparing, sorting, distinguishing), identity itself is relational.

The emerging consensus in analytic metaphysics treats identity as explicable rather than primitive—precisely the view defended here.

3.8 Wiggins: Sortals and Absolute Identity

Wiggins (2001) maintains that identity is absolute (contra Geach) but sortal-governed. When we ask whether $a = b$, we need a sortal F to ground the question: is a the same F as b?

Wiggins treats sortals as providing *criteria* for identity without making identity sortal-relative. But this raises a question: if identity *requires* sortals to be applied, and sortals relate particulars to kinds, is identity not already implicitly relational? I argue that Wiggins' framework, taken seriously, supports the relational thesis: what he treats as application-conditions are better understood as constitutive.

4 Identity in Set-Theoretic Foundations

The philosophical literature reveals a persistent tension: identity is treated as primitive, yet every substantive account of what identity *is* appeals to relational structure. The next three sections examine how this tension plays out in foundational mathematics, beginning with set theory.

4.1 Extensionality Already Makes Identity Relational

In ZFC, sets are identical if and only if they have the same members:

$$\forall x \forall y (x = y \iff \forall z (z \in x \iff z \in y))$$

This is the **Axiom of Extensionality**. Notice what it says: extensionality provides a *relational criterion* for identity. Two sets are the same set precisely when they stand in the same membership relations. The axiom does not "define" identity in a reductive sense—first-order logic still treats identity as primitive—but it provides the criterion by which identity is determined for sets.

This is already a relational account of identity. For any sets A and B:

$$A = B \iff \forall z (z \in A \iff z \in B)$$

Identity is not an intrinsic property that sets possess independently of their relations. Identity is *constituted* by membership relations.

One might object: but the *elements* still have primitive identity. This is true for urelements (if admitted) but not for the pure sets of ZFC, where every set is built from the empty set. And the empty set itself illustrates the relational character of identity.

4.2 The Empty Set is Relationally Defined

The Axiom of Empty Set asserts:

$$\exists x \forall y (y \notin x)$$

The empty set is defined *negatively*: it is the set with no members. But this negative definition presupposes the concept of membership. The empty set is distinguished from non-empty sets precisely by lacking what they have.

Proposition 4.1 (Empty Set is Relationally Defined)**.** *The empty set $\emptyset$ satisfies $R(\emptyset) \neq \emptyset$. That is, there exists at least one entity from which $\emptyset$ is distinguishable.*

Proof. The definition of $\emptyset$ is: $\emptyset$ is the unique set S such that $\forall x (x \notin S)$. Consider any non-empty set—for instance, $\{\emptyset\}$, whose existence is guaranteed by the Pairing axiom. The empty set is distinguishable from $\{\emptyset\}$ by the property that $\emptyset$ has no members while $\{\emptyset\}$ has one. Thus $\delta(\emptyset, \{\emptyset\})$ holds, establishing that $R(\emptyset)$ is non-empty.

We need not (and should not) attempt to *collect* all entities from which $\emptyset$ is distinguishable into a set. In ZFC, the "collection" of all non-empty sets would be a proper class, not a set. But the δ-framework does not require $R(A)$ to be a set in the ZFC sense—only that $R(A) \neq \emptyset$, which requires merely exhibiting a witness. In foundational settings where proper classes are first-class citizens (NBG, MK) or where size issues do not arise (HoTT with universe polymorphism), $R(A)$ can be given a more robust interpretation. The key claim—that $\emptyset$ is defined by its distinguishability from something—holds regardless. ■

4.3 ZFC's Design Choices

A referee has noted that ZFC avoids the universal set and that complements are always relative to a given set. This is correct. But these are *design choices*—motivated by the need to avoid Russell's paradox—not forced by logic itself.

The point is not that ZFC is "circular" (a charge I retract from earlier formulations) but that ZFC's design choices *obscure* the relational character of identity while not eliminating it. Consider:

1. **No universal set**: ZFC avoids U to prevent $\{x : x \notin x\}$ from being well-formed. This is a restriction on comprehension, not a vindication of primitive identity.
2. **Relative complements**: The complement of A is always relative to some $B \supseteq A$. This means "what A is not" is always contextual—*reinforcing* rather than undermining the relational thesis.
3. **Extensionality as criterion**: Set identity is *defined* via membership. The primitiveness of identity in ZFC applies only to urelements (if admitted) and to the membership relation itself.

The upshot: ZFC does not refute the relational thesis. If anything, extensionality *instantiates* it for the domain of sets. The question is whether this relational character generalizes.

5 Category Theory and Structural Identity

If extensionality already makes set-theoretic identity relational, category theory takes the next step: it makes relational characterization the *primary* mode of mathematical description. Where set theory asks "what are the members?", category theory asks "how does this object relate to others?" This shift reveals identity as fundamentally structural—not as a philosophical interpretation imposed on the mathematics, but as a consequence of the categorical method itself.

5.1 Lawvere's Structural Set Theory

In 1964, F. William Lawvere proposed the Elementary Theory of the Category of Sets (ETCS), which axiomatizes set theory categorically rather than membership-theoretically (Lawvere, 1964). In ETCS, sets are characterized not by their elements but by their *universal properties*—how they relate to other sets via morphisms.

As Lawvere later reflected: "What my program discarded was instead the idea of elementhood as a primitive." This is a direct instantiation of the relational thesis: rather than defining sets by intrinsic membership facts, ETCS defines them by their categorical role. A set is what it does, not what it contains.

The empty set in ETCS is the *initial object*—the unique (up to unique isomorphism) object with exactly one morphism to every other object. The singleton is the *terminal object*—the unique object with exactly one morphism from every other object. These characterizations are purely relational; they say nothing about "internal structure" and everything about external relations.

5.2 The Yoneda Lemma

The Yoneda lemma makes the relational thesis mathematically precise (Lane, 1998, Ch. III). An object A in a category is *completely determined, up to isomorphism,* by its morphisms to and from other objects: the representable functor $\mathrm{Hom}(A, -)$ that sends each object B to the set of morphisms from A to B fully characterizes A. Any two objects with the same relational profile—the same pattern of morphisms to all other objects—are isomorphic.

This is the relational thesis in mathematical form. There is no "intrinsic nature" of A beyond its morphism structure. What A *is* is exhausted by how A relates to everything else in the category.

Marquis (2013) develops the philosophical implications: category theory offers "foundations without foundationalism." Rather than building mathematics on primitive elements with primitive identity, categorical foundations build on morphisms (relations) as primary. Objects are, in a precise sense, secondary—they are nodes in a web of relations.

5.3 The Principle of Equivalence

In categorical practice, properties and constructions should be *isomorphism-invariant*: if $A \cong B$, then any categorically meaningful property of A holds of B and vice versa. This principle is so fundamental that violations are informally termed "evil" in the category theory community.

The principle of equivalence embodies a methodological commitment to relational identity. To ask whether two isomorphic objects are "really the same" is, categorically speaking, a malformed question. They are equivalent, and equivalence is all the sameness that matters. The intuition that there might be some "residual difference" between isomorphic objects reflects set-theoretic habits, not categorical reality.

This connects to Awodey (1996)'s argument that category theory instantiates mathematical structuralism: mathematical objects simply *are* positions in structures, characterized by their structural relations. The principle of equivalence is structuralism operationalized. Makkai (1998) formalises this commitment through FOLDS (First Order Logic with Dependent Sorts), a logical framework that restricts the use of equality to contexts where it is invariant under the relevant notion of equivalence. FOLDS captures the categorical norm against "evil" at the level of syntax: one simply cannot formulate non-isomorphism-invariant statements. This is a direct formal precursor to the treatment of identity in HoTT, where the univalence axiom achieves a similar restriction by making identity equivalent to equivalence.

5.4 Higher Categories and the Richness of Identity

The development of higher category theory, particularly ∞-categories, reveals that identity has *rich internal structure* (Lurie, 2009). In an ∞-category:

- Objects are related by morphisms (1-cells).
- Morphisms are related by 2-morphisms (transformations between morphisms).
- 2-morphisms are related by 3-morphisms, and so on.
- This hierarchy continues to infinity.

At each level, "sameness" becomes equivalence at that level: isomorphism of objects, natural isomorphism of functors, equivalence of 2-morphisms, and so forth. The result is that identity is not a simple binary relation but a structure with infinitely many layers.

This undermines any conception of identity as primitive. Primitive identity would be structureless—a bare fact of sameness. But categorical identity has internal structure all the way down (or up). The identity relation between two objects in an ∞-category is itself an ∞-groupoid, which can be arbitrarily complex.

5.5 Categorical Identity Supports the Thesis

Category theory supports the relational thesis at multiple levels:

1. **ETCS**: Sets can be characterized without primitive membership, via universal properties.
2. **Yoneda**: Objects are completely determined by their morphisms—relational data exhausts identity.

3. **Equivalence**: Isomorphism-invariance treats equivalence as the natural notion of sameness.
4. **Higher categories**: Identity has rich internal structure, incompatible with primitiveness.

Homotopy Type Theory synthesizes these categorical insights into a foundational framework where, for types in a universe, identity *just is* equivalence.

6 Homotopy Type Theory: Identity as Path-Structure

Homotopy Type Theory (HoTT) and Univalent Foundations provide a rigorous framework in which identity is *explicitly* relational (The Univalent Foundations Program, 2013; Rijke, 2022). HoTT builds on Martin-Löf Type Theory (MLTT), which introduced identity types as a primitive type-former (Martin-Löf, 1975, 1984). In MLTT, identity between terms $a, b : A$ is not a binary predicate but a *type* $\mathrm{Id}_A(a, b)$, whose inhabitants are witnesses of identification. This type-theoretic treatment was originally understood as encoding a decidable, proof-irrelevant notion of identity. The discovery that it need not be—that identity types can carry rich structure—opened the path from type theory to homotopy theory and ultimately to the framework discussed here. This section argues that HoTT vindicates the relational identity thesis.

6.1 Two Levels of Identity

HoTT distinguishes two related but distinct notions of identity:

- **Identity of elements**: For terms $a, b : A$, the identity type $\mathrm{Id}_A(a, b)$ captures when a and b are identical as elements of A.
- **Identity of types**: For types $A, B : \mathscr{U}$ in a universe $\mathscr{U}$, the identity type $\mathrm{Id}_{\mathscr{U}}(A, B)$ captures when A and B are identical as types.

This distinction matters because the univalence axiom (discussed below) specifically concerns identity of types within a universe, not identity of elements within a type. The two notions interact but should not be conflated.

6.2 Identity Types as Paths

In HoTT, identity between terms a and b of type A is captured by the **identity type** $\mathrm{Id}_A(a, b)$ (The Univalent Foundations Program, 2013, Ch. 1–2). Syntactically, identity types are primitive: Martin-Löf type theory introduces them via formation, introduction, and elimination rules without appeal to paths. For two decades after Martin-Löf's original formulation, it was widely assumed that identity types satisfy the *uniqueness of identity proofs* (UIP) principle: any two inhabitants of $\mathrm{Id}_A(a, b)$ are themselves identical. Hofmann and Streicher (1998) refuted this assumption by constructing a groupoid model in which identity types have genuinely non-trivial structure—distinct identity proofs between the same terms that are not themselves provably equal. This result was decisive: it showed that the identity type is not merely a proposition (true or false) but a space that can be arbitrarily complex. The groupoid model opened the door to viewing identity types as path spaces, and thereby to the entire homotopy-theoretic interpretation of type theory (see Pelayo and Warren, 2014, for a survey aimed at the mathematical community).

The homotopy-theoretic *interpretation* views identity types as spaces of paths from a to b in the space A (see Klev, 2022, for a philosophical survey).

- Types are *interpreted* as spaces (homotopy types).
- Terms are *interpreted* as points in those spaces.

- Identity proofs are *interpreted* as paths between points.
- Higher identity proofs are homotopies—paths between paths.

This path interpretation is model-theoretic: it tells us what identity types *denote* in the simplicial set or ∞-groupoid models, not what they *are* syntactically. Nevertheless, the interpretation is not arbitrary—it is *sound*, meaning anything provable in the syntax holds in the models. Riehl et al. (2022) provide the most current and comprehensive treatment of this soundness, surveying how Voevodsky's original simplicial model has been vastly generalized by Shulman to yield an interpretation of HoTT with strict univalent universes in any ∞-topos. The philosophical import is that the relational structure of identity types is not an artifact of one particular model but holds across the entire landscape of ∞-topos semantics—the richness of identity is a robust feature of the framework, not a peculiarity of simplicial sets. The philosophical significance lies in what this soundness reveals: a formal system with syntactically primitive identity types admits models where identity has rich relational structure.

An identity proof is not a binary truth-value but a *witness*—a constructive term showing *how* a and b are identified. On the path interpretation, this witness specifies a relational route from a to b, not an intrinsic property of "sameness."

6.3 The Philosophy of Identity Types

Ladyman and Presnell (2015) and Ladyman and Presnell (2016) provide a comprehensive philosophical analysis of identity in HoTT. They argue that the identity type's distinctive features—particularly path induction—can be justified on pre-mathematical philosophical grounds, not merely as formal conveniences.

Their central insight is that, under the homotopy interpretation, identity in HoTT is *constituted* by paths, not merely represented by them. When we prove $a = b$ by exhibiting a term of type $\mathrm{Id}_A(a,b)$, we are not discovering a pre-existing identity fact; we are *constructing* the identity through the witness itself. Multiple distinct terms of the identity type (when they exist) represent genuinely different ways of identifying a and b—the identity type has internal structure.

This supports the relational thesis directly, at least within the homotopy-theoretic interpretation. If identity were primitive in the sense of structureless, there would be exactly one "sameness fact" between identical objects. But in HoTT, the identity type $\mathrm{Id}_A(a,b)$ can be inhabited by multiple non-equal terms (in higher types). Identity is not a binary yes/no question but a *type* of identifications, each representing—on the path interpretation—a different relational route. This rich structure is incompatible with primitive, structureless identity.

Shulman (2018) develops these philosophical implications from the perspective of a key architect of the framework, arguing that HoTT encodes a "synthetic" approach to higher equalities in which identity is not an external relation imposed on objects but an intrinsic feature of the spaces they inhabit. On Shulman's account, the richness of identity types is not an artifact of over-engineering but a faithful reflection of mathematical practice, where we routinely work with objects that can be "the same" in multiple non-trivially distinct ways.

Chen (2024) connects these insights to ontic structural realism and philosophy of physics, arguing that univalent foundations provide the mathematical framework for a thoroughgoing structuralism about mathematical and physical objects. The HoTT treatment of identity is not merely a technical convenience but a philosophical stance: identity is constituted by structure.

6.4 The Univalence Axiom

The univalence axiom, due to Vladimir Voevodsky, concerns types within a fixed universe $\mathcal{U}$. For types $A, B : \mathcal{U}$, it states:

$$(A =_{\mathcal{U}} B) \simeq (A \simeq B)$$

That is: for types in a universe, the type of identities between A and B is equivalent to the type of equivalences between them. Two types (in a universe) are identical precisely when they are *structurally equivalent*.

The qualification "in a universe" matters: univalence governs identity of types *qua* elements of a universe type. This is not a limitation but a precise specification of scope. For practical purposes in HoTT, types always live in some universe, so univalence applies broadly.

A technical note on scope: structure identity principles like univalence require certain hypotheses to hold. Specifically, the univalence axiom applies to types within a fixed universe $\mathcal{U}$; extending it across universe levels requires universe polymorphism or explicit transport. Furthermore, deriving that specific mathematical structures (groups, rings, categories) satisfy the univalence principle for their domains requires encoding those structures via appropriate record types and verifying that the induced notion of equivalence matches the mathematical one. These conditions are typically satisfied in practice but should not be elided in foundational discussions.

As Awodey (2014) puts it: "Univalence embodies mathematical structuralism." The axiom does not merely say equivalent things *should be treated as* identical; it says they *are* identical, in the foundational sense. Ladyman and Presnell (2019b) subject this claim to careful scrutiny, extending their Types-as-Concepts interpretation to universes and the univalence axiom specifically. They identify subtleties in Awodey's invariance argument—particularly regarding the role of universe levels—while ultimately endorsing the view that univalence provides a principled characterisation of type identity. Their analysis shows that the philosophical import of univalence is not undermined by its technical complexity but rather deepened by it: the identity of types is a structured affair, mediated by the universe hierarchy.

Angere (2017) examines what kind of identity concept univalent foundations actually encode, drawing on the Carnap-Quine debate about intensionality. Angere argues that UF occupies a distinctive position in the extensional-intensional spectrum: the identity of types is extensional (determined by equivalence) while the identity of terms within a type retains intensional structure (the identity type can have multiple inhabitants). This analysis complicates any simple reading of univalence as "identity is just equivalence" and supports the relational thesis in a nuanced way: identity in UF is not merely structural but *stratified*, with different levels of structure at different levels of the type hierarchy.

Ahrens et al. (2021) generalize this insight into *The Univalence Principle*, proving that in any univalent foundation, equivalent mathematical structures are *indistinguishable*—there is no property that can distinguish them. This is a theorem, not a philosophical interpretation: structural equivalence exhausts identity. The Univalence Principle demonstrates that the relational treatment of identity is not merely consistent but *forced* by the univalent framework.

This is the strongest possible formulation of relational identity. Identity just *is* structural equivalence. There is no residual "primitive identity" over and above the relational structure. Tsementzis (2017) argues that univalent foundations thereby constitute the first mathematically rigorous implementation of philosophical structuralism.

6.5 HoTT Vindicates the Identity Thesis

The development of HoTT shows that:

1. **Identity can be formalized relationally.** The identity type $\mathrm{Id}_A(a,b)$ treats identity as constituted by paths—relational structure—rather than an intrinsic property.
2. **Univalence makes structuralism foundational.** The univalence axiom identifies equivalence with equality, embedding a relational conception of identity into the foundations.
3. **Higher identity structure emerges.** Paths can have paths between them (homotopies), revealing that identity has internal relational structure invisible on the classical view.
4. **This is not a philosophical preference but a mathematical discovery.** HoTT is consistent and has been verified in proof assistants. The relational treatment of identity *works*.

The constructive viability of relational identity in HoTT/UF is further demonstrated by Gratzer et al. (2024), who organize the type of iterative sets V^0 into a locally cartesian closed category with the structure necessary to model extensional type theory internally within HoTT/UF. Their construction—fully formalized in Agda—shows that univalent foundations can accommodate a strict (non-univalent) universe of sets satisfying classical properties, without relying on higher inductive types or other complex extensions. For the identity thesis, this is significant: even within a framework where identity is fundamentally relational (via univalence), one can recover a well-behaved category of sets where identity behaves classically. The classical behavior is not primitive but *derived*—it emerges from the richer relational structure rather than being presupposed. This exemplifies the pattern traced throughout this paper: apparent primitive identity is always recoverable as a special case of relational identity, never the reverse.

The physical significance of this framework is demonstrated by Weatherall (2018), who uses categorical equivalence to argue that the hole argument in general relativity is blocked by standard mathematical practice: diffeomorphism-related spacetime models are "the same" in the categorically relevant sense, so permuting the points of a manifold does not generate a distinct physical possibility. Ladyman and Presnell (2019a) extend this argument into HoTT specifically, showing that in univalent foundations, diffeomorphism-related models are not merely equivalent but *identical*—the identity type between them is inhabited. This is the relational identity thesis applied to physics: the identity of spacetime models is constituted by their structural equivalence, not by some primitive fact about the "points" composing them. That HoTT's treatment of identity resolves a foundational problem in general relativity—not merely in pure mathematics—suggests that the relational character of identity is a feature of identity as such, not a peculiarity of mathematical foundations.

One might object that HoTT is just one foundational system among many. True, but HoTT resolves long-standing problems (like Benacerraf's problem about the nature of mathematical objects) that plagued set-theoretic foundations. Ladyman and Presnell (2018) provide a detailed assessment of whether HoTT provides an adequate foundation for mathematics, concluding affirmatively while identifying areas where further philosophical work is needed. It represents the cutting edge of foundational mathematics (see Corfield, 2020, for a comprehensive philosophical treatment), and its treatment of identity as relational suggests this is the direction of travel.

7 The Identity Problem in Mathematical Structuralism

The preceding sections traced a progressive shift in foundational mathematics from treating identity as primitive toward treating it as structurally constituted. Mathematical structuralism provides a natural test case for the relational thesis, because it makes explicit what set theory and category theory leave implicit: that mathematical objects are positions in structures rather than independently existing entities. This commitment raises a sharp question—how can structuralism account for the identity and distinctness of objects?—whose resolution turns out to require precisely the relational conception of identity defended here.

7.1 Benacerraf's Problem

Benacerraf (1965) posed a foundational puzzle: natural numbers can be identified with multiple, mutually incompatible set-theoretic reductions. In von Neumann's construction, $2 = \{\emptyset, \{\emptyset\}\}$; in Zermelo's, $2 = \{\{\emptyset\}\}$. Both work equally well for arithmetic, yet they are different sets.

Benacerraf's conclusion: numbers are not "really" any particular sets. They are, at best, *positions in the natural number structure*. What matters is that something plays the "2-role"—succeeds 1, precedes 3, is even—not what intrinsic properties it has.

This is already a relational conception of identity. The number 2 is not defined by what it *is* but by where it *sits* in the structure. Identity is positional, not intrinsic. As Frege (1884) had already recognised in his Julius Caesar problem, the question of cross-sortal identity—whether a number could be identical to a non-mathematical object—cannot be settled without specifying a domain and the individuation conditions governing it. Benacerraf's problem is, in this light, a generalisation of Frege's: if numbers are not intrinsically any particular kind of thing, then what fixes their identity? The answer, on the relational view, is structural position.

Beni (2020) complicates this narrative by arguing that Benacerraf's problem poses a genuine challenge for ontic structural realism: the gap between mathematical structures (which exhibit the problem) and physical structures (which OSR treats as fundamental) may be wider than structuralists suppose. If mathematical and physical structure differ in kind, the relational treatment of identity in one domain does not automatically transfer to the other. This is a serious objection, though one that the convergence argument developed below—drawing on independent evidence from category theory, HoTT, and analytic metaphysics—is designed to address.

7.2 The Identity Problem

Keranen (2001) pressed a sharp objection to structuralism. If mathematical objects are mere positions in structures, how can structuralism distinguish structurally indiscernible objects?

Consider the complex numbers i and $-i$. They are *structurally indiscernible*: any structural property true of one is true of the other. They are both square roots of -1, related symmetrically to all other complex numbers. If identity is purely structural, what distinguishes them?

Leitgeb and Ladyman (2008) argue that the identity or distinctness of places in a structure is not to be accounted for by anything external to the structure itself: a structure's identity conditions are *internal*, determined by its automorphism group. On their analysis, i and $-i$ are distinct because the complex field admits a non-trivial automorphism (conjugation) that swaps them—and the existence of such an automorphism is itself a structural fact. This graph-theoretic approach suggests that the identity problem arises from demanding that identity be grounded in *properties* of positions, when it should be grounded in the *symmetry structure* of the whole.

Nevertheless, the challenge remains sharp. This is the *identity problem for realist structuralism*. The structuralist seems forced to either:

1. Deny that $i \neq -i$ (absurd—they are demonstrably distinct);
2. Admit non-structural facts about mathematical objects (abandoning structuralism);
3. Find some structural ground for their distinctness.

7.3 Proposed Responses

Several responses have been offered:

Weak Discernibility (Ladyman, 2005): Objects can be discerned through *irreflexive relations*. If there is a relation R such that $R(i,-i)$ but not $R(i,i)$, then i and $-i$ are weakly discernible. For complex numbers, consider: $i+(-i)=0$ while $i+i\neq 0$. This provides a structural, if indirect, ground for distinctness. Ladyman et al. (2012) provide the definitive formal taxonomy, distinguishing absolute, relative, and weak discernibility and showing how the philosophical question of identity factors into formally separable components. Their analysis reveals that weak discernibility is not a concession to the objector but a principled structural relation: objects are weakly discerned when they stand in an irreflexive relation, and this is a genuine form of individuation. Muller and Saunders (2008) demonstrate that this is not merely a philosopher's construction: fermions in quantum mechanics are weakly discerned by physically meaningful relations (specifically, by anti-correlated spin states), providing a concrete physical instantiation of relational identity.

Primitive Identity (Shapiro, 2008): Shapiro (1997) initially held that structuralism can invoke primitive identity for positions. Objects are identical or distinct as a brute fact, irreducible to structural properties. This preserves structuralism about the *nature* of mathematical objects while conceding that *identity* is not reducible.

Model-Theoretic Individuation (Ketland, 2006): Ketland proposes identifying positions in a structure with their types in the expanded structure—the complete set of properties and relations a position satisfies. This provides a model-theoretic criterion that is both structural and fine-grained enough to distinguish i from $-i$ in contexts where the background theory permits it.

Hybrid Structuralism (Button, 2006): Button argues for distinguishing basic structures (those whose positions are all discernible) from constructed structures (built by permuting or identifying positions). On this view, the identity problem arises only for constructed structures, and the structuralist can maintain that basic structures—which include most of the structures mathematicians actually care about—individuate their positions structurally.

HoTT Resolution (Awodey, 2014): In univalent foundations, the identity problem does not arise in its classical form. Identity is structural equivalence, full stop. The question "are i and $-i$ the same?" is answered by asking whether there is a path between them in the relevant type. There is not: complex conjugation swaps them, but this is an automorphism, not an identity path.

Bundle Theory (Assadian, 2024): Assadian offers an account of structuralist mathematical objects as mereological bundles of structural properties—no substrata, only properties-in-relation. On this view, a number just *is* the bundle of its structural role-properties (being even, succeeding 1, preceding 3), and these properties are possessed essentially. This converges with the relational identity thesis from a direction independent of HoTT: if mathematical objects are bundles rather than bare particulars instantiating properties, then identity is constituted by the relational configuration of the bundle's parts rather than by any primitive "thisness" underlying them. Where the HoTT resolution dissolves the identity problem by making identity equivalent to structural equivalence, Assadian's bundle-theoretic resolution dissolves it by eliminating the substrata that would need to be identified in the first place. Both routes arrive at the same destination: identity is irreducibly relational.

7.4 HoTT Dissolves the Problem

The HoTT resolution is not merely another response but a *dissolution* of the problem. The identity problem arises from a gap between "structural sameness" and "numerical identity." HoTT closes this gap by definition: identity *just is* structural equivalence. There is no "residual" identity question once equivalence is settled.

This completes an arc in the philosophy of mathematics:

1. **Early structuralism**: Objects are positions, but identity remains primitive.
2. **Identity problem**: Exposes tension in combining structural objects with primitive identity.
3. **HoTT resolution**: Identity = equivalence; no gap remains for the problem to exploit.

The relational thesis predicts this trajectory. If identity is constituted by relational structure, then any framework that treats identity as primitive while making objects structural will face tensions. HoTT resolves these tensions by making identity relational at the foundational level.

8 Contemporary Metaphysics of Identity

The argument so far has been largely internal to the philosophy of mathematics and logic. But the relational thesis does not depend on accepting any particular foundational programme. This section shows that contemporary analytic metaphysics—working from independent motivations in the philosophy of science, the metaphysics of grounding, and debates about persistence—has converged on strikingly similar conclusions about the derivative character of identity.

8.1 Identity Criteria as Ground

Fine (2016) argues that identity criteria are *statements of ground*—they say what makes it the case that two things are identical. To give identity conditions for Fs is not merely to describe when Fs are the same, but to explain *why* they are the same.

This has a crucial implication: if identity has grounds, identity is not primitive. The grounding framework in contemporary metaphysics (Correia and Skiles, 2019) treats explanatory relations as fundamental. If identity is grounded in other facts (about essence, structure, or relations), then identity is derivative.

What grounds identity? The candidates Fine considers—qualitative properties, individual essences, structural roles—are all relational in the relevant sense. Identity is grounded in how things relate to other things (including properties, kinds, and structures), not in a primitive "sameness" fact.

8.2 Against Primitive Haecceities

Koslicki (2020) systematically evaluates proposals for cross-world identity criteria. She considers haecceities (primitive "thisnesses"), individual essences, qualitative properties, and structural/formal features.

Koslicki argues against haecceities on multiple grounds: they are explanatorily vacuous, multiply ontological commitments unnecessarily, and conflict with our best scientific understanding. Instead, she defends *individual forms*—structural/formal features that ground identity without invoking primitive thisness.

This is a vote for relational identity from within mainstream analytic metaphysics. If individual forms—essentially structural features—ground identity, then identity is constituted relationally.

8.3 Ontic Structural Realism

French and Ladyman (2003) develop *ontic structural realism* (OSR) as a response to challenges from physics, particularly quantum mechanics. OSR holds that structure is ontologically fundamental; objects (if they exist at all) have only "contextual identity" emerging from their structural positions.

The moderate version maintains objects as metaphysically "thin"—real but not fundamental. Esfeld and Lam (2008) develop this moderate position for the case of spacetime, arguing that objects and relations are on the same ontological footing: neither is reducible to the other, but relations provide the identity conditions for objects. On their account, spacetime points have no identity apart from the metrical and topological relations in which they stand—a direct application of the relational identity thesis to physics. The radical version eliminates objects entirely in favor of pure structural relations. Both versions treat identity as derivative: objects are individuated by their structural role, not by primitive identity.

French and Krause (2006) extend this analysis to quantum particles, arguing that the identity and individuality of particles cannot be understood on classical models. Particles may lack primitive identity altogether, or their identity may be "transcendentally constituted" by structural position.

8.4 Identity as Invariance

Simons (2000) offers a distinctive account: continuants (persisting objects) are "invariants among occurrents under equivalence relations." What persists through time is what remains *invariant* under transformations.

This is explicitly a relational/structural account. Identity through time is not a primitive sameness but a structural fact: the object at t_1 and the object at t_2 are identical because they are related by an appropriate equivalence-preserving transformation. Identity is invariance under equivalence.

Sider (2020) explores "structuralism about individuals"—the view that individuals are secondary, patterns of relations primary. On this view, the fundamental level contains only relational structure; objects are derived from structural nodes. Sider considers this a live metaphysical option, noting that much of our best physics points in a structural direction.

8.5 Convergence in Metaphysics

The convergence is striking. From different starting points and methodological traditions, contemporary metaphysics arrives at similar conclusions:

- **Grounding**: Identity criteria are grounding claims; identity is derivative (Fine).
- **Anti-haecceitism**: Primitive thisness is rejected; structural forms individuate (Koslicki).
- **OSR**: Structure is fundamental; objects have contextual identity (French, Ladyman).
- **Invariance**: Identity is invariance under transformation (Simons).
- **Structuralism**: Individuals are secondary to relational structure (Sider).

This convergence provides abductive support for the relational thesis. When multiple independent research programs reach similar conclusions, the most plausible explanation is that they have tracked something real. The relational character of identity is not a philosophical prejudice but a discovery.

9 Objections and Replies

9.1 Objection: Identity Remains Primitive as a Logical Symbol

"You've shown semantic preconditions for meaningful identity claims. But this doesn't make the identity predicate itself non-primitive—'=' remains in the primitive vocabulary of first-order logic."

Reply: This objection conflates syntactic primitiveness with conceptual independence. The symbol '=' appears in the primitive vocabulary—granted. But syntactic primitiveness means only that the symbol is not *defined via other symbols* in the language. It says nothing about whether the *relation* the symbol expresses is ontologically independent of its relata.

Consider: '$\in$' is also primitive in ZFC, yet no one infers that membership is ontologically independent of sets. The primitiveness of '$\in$' reflects an expressive choice about what to take as undefined, not a metaphysical claim about the self-sufficiency of membership. The same holds for '='. Identity is a relation between terms; without relata, there is nothing for '=' to express.

The paper does not deny that '=' is syntactically primitive. It argues that the *relation* denoted by '=' presupposes referential structure: for $A = B$ to express anything, A and B must be defined, and definition requires distinction. Syntactic primitiveness and conceptual derivativeness are perfectly compatible.

9.2 Objection: This Conflates Semantic and Syntactic Levels

"You move between claims about what terms mean (semantics) and claims about formal axioms (syntax) without being clear about which level you're operating at."

Reply: The thesis operates primarily at the semantic level: it concerns what is required for identity claims to be meaningful, not just formally derivable. But the semantic claim has syntactic consequences: if identity requires reference, then any formal system that treats identity as primitive is implicitly relying on a background semantic interpretation that supplies the referential structure.

The point is that the standard characterization of identity ("a relation each thing bears to itself and nothing else") is *circular at the semantic level*. This circularity is not a syntactic defect but a diagnostic of the fact that identity is not genuinely primitive.

9.3 Objection: Extensionality Solves the Problem

"In set theory, extensionality defines identity for sets. So identity is already relational. What more do you want?"

Reply: This objection *supports* the thesis rather than refuting it. Extensionality shows that for sets, identity is indeed relational. The question is whether this generalizes beyond sets.

For pure mathematics, HoTT shows it does: univalence generalizes extensionality to all types. For metaphysics, the question is whether the relational character of mathematical identity reflects something about identity as such. I argue it does: the success of relational treatments in foundations suggests that relationality is not a peculiarity of sets but a feature of identity itself.

9.4 Objection: Identity in First-Order Logic is Primitive by Design

"First-order logic treats identity as primitive for good reason: it enables a simple, well-understood framework. Your philosophical complaints don't undermine this."

Reply: The question is not whether treating identity as primitive is *useful* but whether it is *correct*—whether it captures what identity actually is. I grant that primitive identity is a useful idealization for many purposes. But the emergence of HoTT, where identity is explicitly relational and the framework is still well-understood and usable, shows that the primitive treatment is not forced on us.

When a useful idealization conflicts with a more accurate picture, we should note the idealization while preferring the accurate picture for foundational purposes.

9.5 Objection: What About Necessary Beings?

"Couldn't there be a necessary being whose identity is primitive—a being that just IS, without needing definition through contrast?"

Reply: Necessity of existence does not entail primitiveness of identity. Even a necessary being must be distinguished from contingent beings to be identified *as* necessary. The concept of "necessary existence" is defined against "contingent existence." The necessary being, if it exists, is defined by what it is not: not contingent, not dependent, not mutable. These negations constitute its identity.

9.6 Objection: This is Just Epistemology

"You've shown that we need reference to know what A is. But that's about knowledge, not being."

Reply: The thesis is ontological. Consider: what would it mean for A to be self-identical if A is not defined? Self-identity of an undefined entity is not a deep metaphysical truth—it is a meaningless string of symbols.

If one insists that A could be self-identical while being undefined, one must explain what "A" refers to in the claim "$A = A$." If it refers to nothing, the claim says nothing. If it refers to something, that something is defined, and definition requires distinction.

10 Consequences

The relational thesis, if sound, has consequences beyond the philosophy of mathematics. This section traces three: a constraint on the metaphysics of properties, a structural account of persistence through change, and implications for the "hard problems" that arise when transformations fail to preserve identity.

10.1 The Relational Character of Properties

If identity is constituted relationally, then a strong consequence follows for the metaphysics of properties: there can be no properties that entities possess wholly independently of their relations. The reasoning is straightforward. Suppose an entity's identity is constituted by its relational position. Then any property P that individuates or partially determines what the entity *is* must figure in that relational constitution. A property that played no role in any relation—that was completely invisible to the entity's relational profile—would be epiphenomenal with respect to identity: it could vary without affecting what the entity is. But then it would not be a property of that entity in any robust sense.

This consequence is conditional: it depends on accepting the relational thesis developed in the preceding sections. One might resist it by distinguishing identity-constituting relations from other relations, permitting "intrinsic" properties that do not bear on identity. This is coherent, but it requires a principled account of why some properties escape the relational web while others do not. The burden of argument shifts to the defender of intrinsic properties.

Physics might seem to offer counterexamples—mass and charge are standardly called "intrinsic properties." But mass is defined operationally by gravitational and inertial behaviour: relations to other masses and to spacetime geometry. Charge is defined by electromagnetic interactions. What appear to be intrinsic properties are dispositional: capacities for certain relations. The physics case thus supports rather than undermines the conditional. This is consistent with the dispositional essentialism defended by French and Ladyman (2003) and the structuralist metaphysics of Ladyman et al. (2007), both of which treat dispositional properties as fundamentally relational.

10.2 Identity Under Transformation

When identity is relational, a natural question arises: when does transformation preserve identity? If A is defined by its referential set $R(A)$, and a transformation T alters the relational structure, under what conditions does the transformed entity $T(A)$ remain the same entity?

The answer suggested by the δ-framework is that *identity is preserved under transformation if and only if the transformation respects the equivalence relations constituting that identity*. More precisely: if the referential set $R(A)$ encodes the distinguishability relations that constitute A's definition, then a transformation preserves A's identity when it preserves the relevant distinguishability structure—when $\delta(A,x)$ holds before and after the transformation for all x in the relevant domain.

This principle has a natural formulation in each of the foundational settings discussed above. In category theory, it corresponds to the requirement that functors preserve isomorphism classes. In HoTT, it corresponds to transport along paths in the identity type (The Univalent Foundations Program, 2013, Ch. 2.3). In set theory, it corresponds to the requirement that maps preserve extensional structure. The formal development of this principle and its application to scale-relative dynamics is pursued in Farzulla (2025d).

The significance is that "hard problems" in multiple domains—consciousness, personal identity, institutional persistence—may be understood as cases where a transformation fails to preserve the relational structure constituting identity. When we coarse-grain a physical system and find emergent "memory terms," or when we discretize a continuous signal and encounter aliasing, or when we nominalize a process and generate pseudo-entities, the apparent complexity at the transformed level is an artefact of the transformation's failure to respect equivalence structure. The relational thesis reframes these as technical problems of transformation design rather than metaphysical mysteries.

10.3 Against Primitive Selfhood

A final consequence concerns the concept of the self. If identity is irreducibly relational, then the notion of a "self" that exists prior to and independently of all relations is incoherent—not because selves do not exist, but because selfhood is constituted by the web of distinguishability relations in which a subject is embedded. This is consistent with Buddhist *anattā* (no-self) doctrine and with the process-relational metaphysics of Whitehead (1929), but it arrives at these conclusions from formal rather than phenomenological or soteriological premises.

The implication for AI and cognitive science is direct: if selfhood is relational, then any system embedded in a sufficiently rich web of distinguishability relations has the structural preconditions for identity. Whether this suffices for *consciousness* is a separate question—one pursued in Farzulla (2025b), which argues from eliminative monism that consciousness is not a substance but a nominalization of process, consistent with the relational thesis developed here. The relational thesis removes one obstacle to attributing identity to non-biological systems—the assumption that identity requires some primitive, substrate-dependent "thisness" that only biological entities can possess (Farzulla, 2025c). The formal framework for analysing how relational identity interacts with multi-agent coordination—specifically, the friction costs that arise when agents must negotiate identity across different representational frameworks—is developed in Farzulla (2025a).

11 Conclusion

The law of identity is not foundational. $A = A$ presupposes that A is defined, and definition requires reference to what A is not. Identity is therefore derivative of referential structure.

11.1 The Convergence of Four Traditions

This thesis finds support from four independent research programs, each with different methodologies, starting points, and concerns:

Homotopy Type Theory. Identity types are path spaces; identity is constituted by the paths connecting points, not by an intrinsic sameness property. The univalence axiom makes structural equivalence identical to equality—literally, not merely analogically. As Ladyman and Presnell (2016) emphasize, this is not a philosophical interpretation but a mathematical fact.

Category Theory. The Yoneda lemma proves that objects are completely determined by their morphisms. ETCS (Lawvere, 1964) axiomatizes set theory without primitive membership. The principle of equivalence treats isomorphism as the natural notion of sameness. Marquis (2013) draws the conclusion: categorical foundations treat relational structure as primary, objects as secondary.

Mathematical Structuralism. The trajectory from Benacerraf's problem through Keranen's identity problem to the HoTT resolution traces a path toward relational identity. Hellman and Shapiro (2018) survey this landscape, noting that the most promising structuralist foundations are those that treat identity as constituted by structure rather than primitive.

Analytic Metaphysics. Fine's grounding framework, Koslicki's rejection of haecceities, ontic structural realism, and Simons' invariance account all converge on treating identity as derivative. As Sider (2020) notes, structuralism about individuals—the view that relational structure is primary and individuals secondary—is a live metaphysical option supported by our best physics.

This convergence is significant. These traditions developed largely independently, with different motivations:

- HoTT emerged from computer science and homotopy theory.
- Category theory from algebraic topology and universal algebra.
- Structuralism from philosophy of mathematics.
- OSR from philosophy of physics and quantum mechanics.
- Grounding metaphysics from debates about fundamentality.

Yet they reach similar conclusions about identity. When independent inquiries converge, the most plausible explanation is that they have tracked something real. The relational character of identity is not an artifact of any particular methodology but a feature of identity itself.

11.2 The Trajectory of Foundations

The historical trajectory confirms this conclusion. Foundational mathematics has moved steadily from treating identity as primitive toward treating it as constituted by structure:

- **1889**: Peano's axioms treat identity primitively.
- **1908**: Zermelo's axiomatization uses extensionality—relational identity for sets.
- **1964**: Lawvere's ETCS drops primitive membership entirely.
- **2009**: Higher category theory reveals identity has rich internal structure.
- **2013**: HoTT makes univalence foundational—identity *is* equivalence.

This is not random drift but directed progress. Each step makes identity more explicitly relational, because that is what works mathematically. The relational thesis predicted this trajectory; the trajectory confirms the thesis.

11.3 Implications

The relational thesis has both negative and positive implications. Negatively, it constrains what can count as a satisfactory account of identity: any framework that treats identity as primitive and structureless is, at best, an idealization that obscures the relational constitution it depends on. Positively, it offers a unified perspective on problems that appear disparate when identity is taken as primitive.

Consider the cluster of "hard problems" that arise when descriptions change level: memory terms in coarse-grained statistical mechanics, aliasing in discretized signals, pseudo-entities generated by nominalization in natural language. On the relational view, these are not independent puzzles but instances of a single phenomenon: *transformation failure*. When a transformation between descriptive levels fails to preserve the equivalence relations constituting identity at the original level, artefacts emerge at the transformed level that masquerade as genuine structure. The relational thesis predicts this and provides a diagnostic: check whether the transformation respects the distinguishability relations encoded in $R(A)$.

More broadly, the thesis suggests that the traditional boundary between formal and material identity—between the identity of mathematical objects and the identity of physical or social entities—is less sharp than usually assumed. If identity is constituted by relational structure in all these domains, then the formal tools developed in HoTT, category theory, and structuralist metaphysics are not merely analogies for thinking about physical and social identity but genuine theoretical resources.

Identity is irreducibly relational. What remains is to work out the consequences—for transformation theory, for the metaphysics of persistence, and for the design of formal systems adequate to the relational character of the world they aim to describe.

Acknowledgements

Portions of this manuscript were drafted and revised with assistance from Claude (Anthropic). The author retains full responsibility for all intellectual content, analytical decisions, and interpretive claims. This work was conducted as part of the Adversarial Systems Research program at Dissensus AI. Working papers and replication materials are available through the Adversarial Systems & Complexity Research Initiative (ASCRI). Comments and correspondence are welcome at murad@dissensus.ai.

Declarations

Conflict of Interest. The author declares no competing interests.

Funding. This research received no external funding.

AI Assistance. Claude (Anthropic) assisted with manuscript drafting, formal argumentation, and literature synthesis. All research design, analysis, and interpretation are solely the author's.

Author Contributions. Sole author.

References

Benedikt Ahrens, Paige Randall North, Michael Shulman, and Dimitris Tsementzis. The univalence principle. *arXiv preprint arXiv:2102.06275*, 2021.

Staffan Angere. Identity and intensionality in univalent foundations and philosophy. *Synthese*, 198: 1177–1217, 2017. doi: 10.1007/s11229-016-1301-z.

Bahram Assadian. Mathematical structuralism and bundle theory. *Ratio*, 2024. doi: 10.1111/rati.12397. Early View.

Steve Awodey. Structure in mathematics and logic: A categorical perspective. *Philosophia Mathematica*, 4(3):209–237, 1996.

Steve Awodey. Structuralism, invariance, and univalence. *Philosophia Mathematica*, 22(1):1–11, 2014.

Paul Benacerraf. What numbers could not be. *The Philosophical Review*, 74(1):47–73, 1965.

Ermanno Bencivenga. Free logics. In Dov M. Gabbay and Franz Guenthner, editors, *Handbook of Philosophical Logic*, volume 5, pages 147–196. Springer, 2nd edition, 2002.

Majid Davoody Beni. The Benacerraf problem as a challenge for ontic structural realism. *Philosophia Mathematica*, 28(1):35–59, 2020. doi: 10.1093/philmat/nkz022.

Max Black. The identity of indiscernibles. *Mind*, 61(242):153–164, 1952.

Tim Button. Realistic structuralism's identity crisis: A hybrid solution. *Analysis*, 66(3):216–222, 2006.

Massimiliano Carrara and Ciro De Florio. Identity criteria: An epistemic path to conceptual grounding. *Synthese*, 197:3151–3169, 2020. doi: 10.1007/s11229-018-1878-5.

Lu Chen. Univalence and ontic structuralism. *Foundations of Physics*, 54(3):1–27, 2024. doi: 10.1007/s10701-024-00768-4.

David Corfield. *Modal Homotopy Type Theory: The Prospect of a New Logic for Philosophy*. Oxford University Press, 2020.

Fabrice Correia and Alexander Skiles. Grounding, essence, and identity. *Philosophy and Phenomenological Research*, 98(3):642–670, 2019.

Michael Esfeld and Vincent Lam. Moderate structural realism about space-time. *Synthese*, 160(1): 27–46, 2008. doi: 10.1007/s11229-006-9076-2.

Murad Farzulla. The axiom of consent: Friction dynamics in multi-agent coordination. *arXiv preprint arXiv:2601.06692*, 2025a. doi: 10.48550/arXiv.2601.06692.

Murad Farzulla. Replication optimization at scale: Dissolving the hard problem via Occam's razor. *Zenodo Preprint*, 2025b. doi: 10.5281/zenodo.18013187.

Murad Farzulla. Relational functionalism: Friendship as substrate-agnostic process. *Zenodo Preprint*, 2025c. doi: 10.5281/zenodo.17626860.

Murad Farzulla. ROM: Scale-relative formalism for persistence-conditioned dynamics. *arXiv preprint arXiv:2601.06363*, 2025d. doi: 10.48550/arXiv.2601.06363.

Kit Fine. Essence and modality. *Philosophical Perspectives*, 8:1–16, 1994.

Kit Fine. Identity criteria and ground. *Philosophical Studies*, 173:1–19, 2016.

Gottlob Frege. *Die Grundlagen der Arithmetik*. Koebner, Breslau, 1884. Trans. J.L. Austin as *The Foundations of Arithmetic*, Blackwell, 1950.

Steven French and Décio Krause. *Identity in Physics: A Historical, Philosophical, and Formal Analysis*. Oxford University Press, 2006.

Steven French and James Ladyman. Remodelling structural realism: Quantum physics and the metaphysics of structure. *Synthese*, 136:31–56, 2003.

Andre Gallois. Identity over time. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, 2016.

Peter T. Geach. Identity. *The Review of Metaphysics*, 21(1):3–12, 1967.

Daniel Gratzer, Håkon Robbestad Gylterud, Anders Mörtberg, and Elisabeth Stenholm. The category of iterative sets in homotopy type theory and univalent foundations. *Mathematical Structures in Computer Science*, 2024. doi: 10.1017/S0960129524000070.

Geoffrey Hellman and Stewart Shapiro. *Mathematical Structuralism*. Cambridge Elements in the Philosophy of Mathematics. Cambridge University Press, 2018.

David Hilbert and Paul Bernays. *Grundlagen der Mathematik*, volume 1. Springer, 1934. Contains the definition of identity via indiscernibility in higher-order logic.

Martin Hofmann and Thomas Streicher. The groupoid interpretation of type theory. In Giovanni Sambin and Jan M. Smith, editors, *Twenty-Five Years of Constructive Type Theory*, pages 83–111. Oxford University Press, 1998.

Jukka Keranen. The identity problem for realist structuralism. *Philosophia Mathematica*, 9(3):308–330, 2001.

Jeffrey Ketland. Structuralism and the identity of indiscernibles. *Analysis*, 66(4):303–315, 2006. doi: 10.1093/analys/66.4.303.

Ansten Klev. Identity in martin-löf type theory. *Philosophy Compass*, 17(2), 2022. doi: 10.1111/phc3.12805.

Kathrin Koslicki. Essence and identity. In Mircea Dumitru, editor, *Themes from Kit Fine*. Oxford University Press, 2020.

Saul Kripke. *Naming and Necessity*. Harvard University Press, 1980.

James Ladyman. Mathematical structuralism and the identity of indiscernibles. *Analysis*, 65(3):218–221, 2005.

James Ladyman and Stuart Presnell. Identity in homotopy type theory, part i: The justification of path induction. *Philosophia Mathematica*, 23(3):386–406, 2015.

James Ladyman and Stuart Presnell. Identity in homotopy type theory, part ii: The conceptual and philosophical status of identity in hott. *Philosophia Mathematica*, 25(2):210–245, 2016.

James Ladyman and Stuart Presnell. Does homotopy type theory provide a foundation for mathematics? *British Journal for the Philosophy of Science*, 69(2):377–420, 2018. doi: 10.1093/bjps/axw006.

James Ladyman and Stuart Presnell. The hole argument in homotopy type theory. *Foundations of Physics*, 50:319–329, 2019a. doi: 10.1007/s10701-019-00293-9.

James Ladyman and Stuart Presnell. Universes and univalence in homotopy type theory. *The Review of Symbolic Logic*, 12(3):426–455, 2019b. doi: 10.1017/S1755020316000460.

James Ladyman, Don Ross, David Spurrett, and John Collier. *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press, 2007.

James Ladyman, Øystein Linnebo, and Richard Pettigrew. Identity and discernibility in philosophy and logic. *The Review of Symbolic Logic*, 5(1):162–186, 2012. doi: 10.1017/S1755020311000281.

Karel Lambert. *Philosophical Applications of Free Logic*. Oxford University Press, 1991.

Saunders Mac Lane. *Categories for the Working Mathematician*. Springer, 2nd edition, 1998.

F. William Lawvere. An elementary theory of the category of sets. *Proceedings of the National Academy of Sciences*, 52(6):1506–1511, 1964.

Scott Lehmann. More free logic. *Handbook of Philosophical Logic*, 5:197–259, 2002.

Hannes Leitgeb and James Ladyman. Criteria of identity and structuralist ontology. *Philosophia Mathematica*, 16(3):388–396, 2008. doi: 10.1093/philmat/nkm039.

Øystein Linnebo. *Thin Objects: An Abstractionist Account*. Oxford University Press, 2018.

Jacob Lurie. *Higher Topos Theory*, volume 170 of *Annals of Mathematics Studies*. Princeton University Press, 2009.

Michael Makkai. Towards a categorical foundation of mathematics. In *Logic Colloquium '95*, volume 11 of *Lecture Notes in Logic*, pages 153–190. Springer, 1998.

Jean-Pierre Marquis. Categorical foundations of mathematics or how to provide foundations for abstract mathematics. *Review of Symbolic Logic*, 6(1):51–75, 2013.

Per Martin-Löf. An intuitionistic theory of types: Predicative part. In *Logic Colloquium '73*, pages 73–118. North-Holland, 1975.

Per Martin-Löf. *Intuitionistic Type Theory*. Bibliopolis, Naples, 1984. Notes by Giovanni Sambin.

F. A. Muller and Simon Saunders. Discerning fermions. *British Journal for the Philosophy of Science*, 59(3):499–548, 2008. doi: 10.1093/bjps/axn027.

John Nolt. Free logic. *The Stanford Encyclopedia of Philosophy*, 2020. Fall 2020 Edition.

Álvaro Pelayo and Michael A. Warren. Homotopy type theory and Voevodsky's univalent foundations. *Bulletin of the American Mathematical Society*, 51(4):597–648, 2014. doi: 10.1090/S0273-0979-2014-01456-9.

Willard Van Orman Quine. Ontological relativity. *Journal of Philosophy*, 65(7):185–212, 1968.

Willard Van Orman Quine. *Theories and Things*. Harvard University Press, 1981.

Willard Van Orman Quine. *Philosophy of Logic*. Harvard University Press, 2nd edition, 1986.

Emily Riehl, Kenji Maillard, David Michael Roberts, Michael Shulman, Jonathan Sterling, and Andrew W. Swan. On the ∞-topos semantics of homotopy type theory. *Bulletin of the London Mathematical Society*, 2022. doi: 10.1112/blms.12754.

Egbert Rijke. *Introduction to Homotopy Type Theory*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2022. arXiv:2212.11082.

Stewart Shapiro. *Philosophy of Mathematics: Structure and Ontology*. Oxford University Press, 1997.

Stewart Shapiro. Identity, indiscernibility, and ante rem structuralism: The tale of i and -i. *Philosophia Mathematica*, 16(3):285–309, 2008.

Michael Shulman. Homotopy type theory: A synthetic approach to higher equalities. In Elaine Landry, editor, *Categories for the Working Philosopher*, pages 36–57. Oxford University Press, 2018.

Theodore Sider. *The Tools of Metaphysics and the Metaphysics of Science*. Oxford University Press, 2020.

Peter Simons. Identity through time and trope bundles. *Topoi*, 19:147–155, 2000.

The Univalent Foundations Program. *Homotopy Type Theory: Univalent Foundations of Mathematics*. Institute for Advanced Study, 2013. URL `https://homotopytypetheory.org/book/`.

Dimitris Tsementzis. Univalent foundations as structuralist foundations. *Synthese*, 194:3583–3617, 2017. doi: 10.1007/s11229-016-1109-x.

James Owen Weatherall. Regarding the 'hole argument'. *British Journal for the Philosophy of Science*, 69(2):329–350, 2018. doi: 10.1093/bjps/axw012.

Alfred North Whitehead. *Process and Reality*. Macmillan, 1929.

David Wiggins. *Sameness and Substance Renewed*. Cambridge University Press, 2001.