ChatGPT

Beyond Classical Identity: A First-Order Theory of Provenance, Continuity, and Constitution

Abstract

This paper develops a general first-order theory of identity that goes beyond the classical treatment of identity as mere sameness by properties. The thesis is that numerical identity should be explained not just by formal logical laws but by deeper structures: **provenance** supplies the basis for numerical identity of concreta and abstracta that come into being, **continuity** governs persistence through time under an explicit anti-branching policy, and **constitution** captures aspectual or constitutive sameness without collapsing it into numerical equality. The framework is a many-sorted first-order calculus that is conservative over standard first-order logic with equality on ordinary domains, preserving the usual meta-properties of identity where they belong. At the same time, it offers solutions to familiar stress tests of identity: cases of duplication, fission and fusion, the Ship of Theseus puzzle, material constitution puzzles like the statue and the clay, and intensional opacity in substitutivity. The theory is presented as a general "identity calculus" useful in metaphysics, law, and information systems. Compatibility with a relational-first **Catholic theology** (for example, modeling Trinitarian claims) is discussed as an application **rather** than as a driving constraint of the theory. (All citations are in Chicago author-date style.)

Introduction

Philosophers have long sought to understand what makes something *numerically one and the same* beyond the bare formal properties of the identity predicate. The classical view treats identity as an *absolute equivalence relation* (reflexive, symmetric, transitive) satisfying Leibniz's Law (the indiscernibility of identicals) [1]. In that formal sense, identity itself is trivial—each thing is identical to itself and to nothing else, as David Lewis famously remarked (Lewis 1986, 192–93) [2]. However, simply insisting on the formal constraints of identity does not explain *why* or *how* a thing is one, especially in philosophically puzzling cases. There is a crucial distinction often conflated in discussions of identity: the **formal constraints** that any identity relation must satisfy versus the **truth-makers or grounds** that make identity claims true in real situations. Classical first-order logic codifies the formal constraints (e.g. $\forall x\, x = x$; $\forall x \forall y\, (x = y \rightarrow (Fx \rightarrow Fy))$ as an axiom schema) [1], but it leaves open what *grounds* the truth of "x is the same as y" in the world.

This paper's central move is to keep the classical constraints in their proper place (ensuring that identity in each domain remains an equivalence relation obeying Leibniz's Law [1]), but to change the **truth-makers** for identity by introducing new explanatory notions. We propose to *anchor* numerical identity in **provenance** (the origin or coming-to-be of an entity), to explicate persistence through time via **continuity** (spatio-temporal or causal continuity under an explicit no-branching policy), and to analyze aspectual sameness or constitutive unity via a **constitution** relation (a sameness-in-aspect that never collapses into outright identity across categories). By relocating the explanatory basis in this way, we can solve long-standing metaphysical puzzles while respecting the formal laws of identity. The payoffs include clean resolutions of diachronic paradoxes and a principled handling of opaque or "typed" contexts where naive substitutivity of identicals leads to incorrect inferences [3]. In other words, the theory preserves what is

right about classical identity (its formal rigor and non-relativized unity) but augments it with additional structure that accounts for what makes something one or the same in practice.

To motivate the need for such a theory, consider some familiar problems. Since antiquity, philosophers from Aristotle through Avicenna and Aquinas have wrestled with how an entity remains the *same* through change or under different descriptions. Aristotle distinguished **numerical unity** from mere qualitative similarity, noting that only one and the same substance can underlie change in qualities (Aristotle 1984) [4]. Medieval thinkers like Avicenna and Aquinas sought principled "individuating factors" (such as existence instantiated in matter, or designated matter) that make an individual this particular individual, beyond its general essence (Avicenna [1027] 2005; Aquinas [1265] 1981). In modern terms, Leibniz formulated the classical laws of identity (the indiscernibility of identicals and identity of indiscernibles) and pointed out that identity is *necessary* if true (Leibniz [1686] 1989). Yet these formal or essentialist insights do not by themselves settle concrete identity puzzles. Kripke (1980), for example, showed that knowing the formal fact that Clark Kent = Superman does not license substituting one name for the other in Lois Lane's belief context without potential falsehood – a substitution that classical logic would naively permit [3]. Likewise, classic puzzles like the Ship of Theseus show that the formal laws alone do not determine *which* future individual is identical to an original when two qualitatively continuous candidates emerge. We need more than the bare "=" symbol and its laws; we need an account of *what makes something itself over time* and *in different respects*.

In what follows, we introduce a new first-order framework for identity that addresses these needs. Section II reviews the background and limits of the classical picture: how identity's formal laws function as meta-requirements and why they underdetermine answers in cases of **directional dependence**, **temporal sequence**, **reconstitution of matter**, and **intensional contexts**. We will see that the goal is not to reject fundamental laws like reflexivity, symmetry or transitivity – those remain non-negotiable – but rather to supplement them with principled truth-makers that avoid paradox. Section III lays out the **language and semantics** of our many-sorted first-order theory, including sort-specific identity and three new primitives: a provenance tag function, a continuity relation, and a guarded constitution operator. Section IV presents the **axioms** governing these primitives, showing that provenance-based identity is an equivalence relation that coincides with classical identity on each sorted domain, that continuity (given an anti-branching policy) yields a one-per-continuant equivalence on world-lines, and that constitution is a constrained, non-transitive sameness-in-aspect relation. Section V sketches the **meta-theory**, proving that our system is a conservative extension of standard first-order logic with equality and retains completeness, compactness, and model existence (indeed, it admits intuitive models both with and without fission phenomena). Section VI works through **case studies**: Ship of Theseus, biological fission/fusion, teletransportation and duplication, the statue-and-clay puzzle, and intensional identity puzzles. We will see how the new theory yields clean, commonsensical solutions in each scenario – for example, by assigning a single origin-tag to Theseus's ship and thereby declaring the reconstructed ship of old planks a numerically distinct vessel [5], or by using the constitution operator to say the statue is "the same *material object* as the clay" without identifying it outright (thus avoiding Leibniz's Law violations regarding their different properties [6]). Section VII compares our approach to **competing views**: we retain the classical absolute identity relation rather than adopting **relative identity** (Geach 1967) [7]; we avoid reliance on Lewisian **counterpart theory** for intra-world identity decisions; and we handle essences not by building them into the identity criteria of objects, but by letting essential properties govern the constitution relation (thus side-stepping certain **essentialist** identity-condition controversies). Section VIII outlines applications beyond pure metaphysics: we show how the calculus can model identity and change in information systems (versioning, deduplication), law (chain of title and continuity of organizations), biology (organism lineage and traceability), and even theology (offering a formal model for the Christian Trinity consistent with orthodox

doctrine, following the spirit of Brower and Rea's (2005) constitution approach). Finally, Section IX considers **objections and replies**. We address concerns that (1) provenance-based identity collapses into a mere assertion of primitive thisness or *haecceity*, (2) that building a policy into identity compromises its objectivity, and (3) that our constitution operator is just a disguised form of relative identity. In each case we argue that the new framework either avoids the issue or turns it into a virtue (for example, by making policies explicit and non-arbitrary, or by enforcing a strict separation between aspectual sameness and true identity). We conclude that relocating the explanatory basis of identity – provenance for numerical unity, continuity for persistence, constitution for aspectual sameness – allows us to respect the classical logic of identity while solving puzzles that have long plagued both philosophers and practitioners.

Before proceeding, it is worth situating this work in the broader context of identity theory. The idea of supplementing the bare logic of "=" with additional structure has precedents: Wiggins (2001) insisted that identity statements are only meaningful relative to a sortal that provides a criterion of identity (the present theory can be seen as fleshing out such criteria in a general logical form). Noonan and Curtis (2018) survey many of the puzzles we target and discuss solutions like *best-candidate* or *anti-branching* rules; our continuity relation can be viewed as a formalization of an anti-branching criterion in first-order semantics. Philosophers such as van Inwagen (1990), Sider (2001), and Hawley (2001) have debated perdurantist vs endurantist accounts of persistence; our approach to continuity is compatible with either view in spirit (one can think of continuity classes as perduring worms or as persistently existing individuals) but importantly introduces a *policy parameter* to handle indeterminacy in branching cases. In handling constitution, we build on insights from Fine (1999) and Thomasson (2007) that one can acknowledge distinct but coincident objects (or object and matter) without contradiction, by carefully distinguishing the relations of unity involved. And in acknowledging vagueness or indeterminacy in identity claims (in a controlled way), we align with the supervaluationist tradition (Fine 1975) that shows one can accommodate borderline cases while preserving classical logic [8]. The novelty here is to integrate these strands into a single coherent first-order theory. The result, we hope, is both formally rigorous and philosophically illuminating.

(Citations in the text are given in author-date format; full references are in the Bibliography.)

Background and Limits of the Classical Picture

Classical first-order logic treats identity as a primitive two-place predicate "=" subject to axioms that ensure it behaves as true equality. In particular, the classical laws require that identity be reflexive (everything is identical to itself), symmetric, and transitive – i.e. an **equivalence relation** [9] – and that it obey **Leibniz's Law**, the indiscernibility of identicals (if $a = b$ then any property true of a is true of b as well) [1]. Together with its contrapositive (if some property distinguishes a and b then $a \neq b$), Leibniz's Law captures our basic understanding that an identity statement has robust consequences: we routinely infer from "Fa and ¬Fb" that a is not b [10]. These formal constraints are treated as *meta-requirements* on any admissible interpretation of "=". They ensure that identity, in Quine's words, is "utterly simple and unproblematic" in itself (Quine 1964, Lucretius to Quine 1980) – or as Lewis put it, it is never a mystery what makes something identical to itself [2].

Yet the **classical picture** by itself underdetermines many important and puzzling cases. The formal laws can be perfectly satisfied while we still lack answers to questions like: *Is this future person the same as that present person?* or *Can two things in one place be one or two?* or *Does a believer's ignorance block substitution?* The reason is that the formal laws are *silent* on anything that goes beyond the logical form of identity statements. They do not tell us which conditions in the world count as sufficient or necessary for identity

over time or across changes. Thus, when we confront cases involving **directional dependence** (e.g. an object and its origin or cause), **temporal sequence** (an object existing at different times or undergoing alteration), **reconstitution** (an object and the matter that makes it up, or an object that is dismantled and rebuilt), or **intensional contexts** (beliefs, modal descriptions, etc. where substituting co-referential terms can change truth value), the classical laws alone do not resolve the truth of identity claims.

For a compact survey, consider a few classic examples:

- **Ship of Theseus (Reconstitution and Fission):** A famous puzzle described by Plutarch (and elaborated by Hobbes) asks us to imagine a ship that over time has all its wooden planks replaced one by one. The original planks are not discarded but stored, and eventually they are reassembled in their original configuration. At the end of the process, we have two ships: the continuously repaired ship (with entirely new planks) and the reassembled ship made of all the original planks. Which one, if either, **is** the original Ship of Theseus? The classical laws of identity guarantee that if either candidate is identical to the original, it must satisfy transitivity and other properties, but they **do not tell us which candidate to choose**. In fact, if one naively claimed both are identical to Theseus's ship, one would violate the transitivity of identity because the two resulting ships are clearly not identical to each other [11] [12]. Classical logic forbids one thing from becoming two, but it doesn't tell us *which one* carries the original identity. Many proposed solutions impose extra criteria: some say the gradually replaced ship "inherits" the identity (it has continuity of form and function), making the reassembled one a mere copy [5]; others argue the opposite; still others say the identity is indeterminate or that we must distinguish "same ship" in different senses [13] [14]. The takeaway is that the formal notion of identity as such is consistent with either choice – additional principles are needed to explain numerical sameness over such a process.

- **Kripke's Puzzles of Belief (Intensional Contexts):** In a well-known example, Kripke (1979) imagines a person who knows of a famous pianist under the name "Paderewski" but is unaware that Paderewski was also a politician. This person might assert "Paderewski has musical talent" while denying "Paderewski has political savvy," because he thinks perhaps there are two different Paderewskis. In reality it's the same man. Formally, if we let a = Paderewski the musician and b = Paderewski the statesman, we have $a = b$. Classical logic would allow substitution of a for b in any context, but in the context of someone's belief or knowledge report, substituting can lead from truth to falsehood. Similarly, the Morning Star = the Evening Star (both Venus), yet a naive subject can believe "The Morning Star is Venus" is true and "The Evening Star is Venus" is false without contradiction in ordinary language. These are cases of **opaque contexts** where co-referential terms are not freely interchangeable salva veritate [3]. The formal law of indiscernibility of identicals technically is not violated (Leibniz's Law applies only to all *actual* properties of the object, not to properties of names or representations [15]). But these puzzles show that a *truth-maker* approach is needed: the identity is true, yet something about the cognitive or linguistic context prevents straightforward substitution. The classical picture itself does not explain this; one must invoke a theory of *how identity interacts with intensional contexts* (for example, by saying the content of belief is sensitive to modes of presentation or that belief contexts are not extensional). In our framework, this will be handled by a **No Mixing** constraint that prevents aspectual or contextual equivalences from licensing substitution into arbitrary contexts.

- **Statue and Clay (Constitution puzzle):** A sculptor forms a statue from a lump of clay. Intuitively, the lump of clay existed before the statue did; the statue came into being at a certain time (when the

sculptor shaped it) and might be destroyed (by smashing it) while the clay still exists (perhaps just deformed). Are the statue and the clay identical, or are they two different co-located things? If we say they are identical, we face immediate problems: the clay has the property "could survive squashing into a ball" whereas the statue does not (the statue would cease to exist if squashed, losing its form). By Leibniz's Law, if there is any property one has and the other lacks, they cannot be strictly identical [6]. So most philosophers conclude the statue ≠ the clay (they are distinct objects that just happen to coincide materially). But then how can two distinct objects share all the same matter at a time and occupy exactly the same space? And which one is the *one* object in that place? Again, the formal laws of identity don't tell us how to handle this; they only forbid calling them identical if they truly have different properties. One approach is to say the statue and clay are identical in some *aspects* (say, they are the same material object or composed of the same atoms at that time) but not identical *simpliciter*. This hints at a relation weaker than identity – often called **constitution** – which the classical picture doesn't have. If one tries to force everything into the classical mold, one gets paradoxes or counter-intuitive proliferations of entities (e.g. saying there are two coincident objects, with no explanation beyond brute fact). We need a way to say "X and Y are the same in aspect A" without saying "X = Y" outright, to solve this puzzle.

- **Duplication, Fission, and Fusion (Identity over Time):** Beyond the Ship of Theseus, modern thought experiments abound in which persons or objects split or merge. For instance, science fiction scenarios of teleportation often imagine that a person's body is disintegrated in one location and perfectly reassembled somewhere else. If the original is destroyed and only one copy is made, many people are inclined to say the person **survived** (teletransportation was just a form of travel). But if the original were not destroyed – if it were an accidental duplicate – now there are two qualitatively identical claimants to be "the same person" as the original. They cannot both be identical to the original by transitivity (and by the fact they are not identical to each other). Which, if either, is the original person? Similarly, in biology, an organism like a flatworm can be cut in half and each half regenerates into a complete worm; which resulting worm (if any) is identical to the original? Conversely, sometimes two organisms can fuse (e.g. in colonial organisms or hypothetical mind mergers) – which identity continues? The classical laws by themselves only tell us that identity is one-to-one and cannot branch. But they do not contain a built-in rule for *resolving* a branching scenario. They also do not allow "partial identity" or temporary identity if one considers the idea that something could be identical up to a point and then not (most philosophers reject the notion of "temporary identity" as incoherent because of Leibniz's Law and the all-or-nothing nature of identity [16]). Hence, solutions have been proposed that *supplement* classical identity with additional conventions or criteria: for example, the **closest continuer** theory (Nozick 1981) says whichever future being has the greatest degree of continuity with the original is the "real" continuation (the other is a copy), effectively imposing a policy to choose one branch. Another approach is to say identity simply doesn't apply in fission cases – the original ceases and two new entities begin, even though psychologically continuous with it (Parfit 1984). In short, the problem is not the formal properties of "=" – those remain intact – but their insufficiency. We must *augment* identity with a notion of **continuity** plus an **anti-branching policy** to have a determinate answer in such cases, or else accept indeterminate identity (which leads to thorny logical issues, as argued by Evans 1978). Our framework will take the route of explicitly building in a continuity relation and a non-branching criterion, such that in a model where branching occurs, identity statements involving the post-branch individuals can be marked as neither true nor false (needing further specification) unless a policy picks a winner.

The above examples highlight the *pressure points* where the classical picture of identity shows its limits. Importantly, none of this invalidates the core laws themselves – we are **not** aiming to deny reflexivity, symmetry, transitivity, or Leibniz's Law. Those are preserved in our theory, but they are kept as high-level constraints (meta-properties) that any adequate identity relation must obey. The **target** of our revision is not the formal properties, but the **grounds of identity claims**. We want to supply a richer story such that when we say "X is (numerically) identical to Y," there is a substantive reason – X and Y share an origin or "tag" in our theory. And when we say "X persists as Y over time," it's because there is a continuous sequence from X to Y under a policy that disallows divergence. And when we say "X is in some sense the same as Y" (statue vs clay, or person vs body, or even Father vs Son in theology), we can formalize that as a **constitution relation** that mimics identity in controlled ways but never tricks us into an illicit substitution in a different context.

In summary, the classical picture treats identity as a simple logical primitive and its laws as boundary conditions. Those laws themselves are **correct** so far as they go; they are constitutive of what we *mean* by identity (e.g. it wouldn't be identity if it weren't transitive, etc.). However, as many authors have noted (Wiggins 2001; Noonan and Curtis 2018; Sider 2001), to solve real identity puzzles we need to **go beyond** the classical picture. We need additional relations or apparatus to handle *continuity over time, composition and constitution, causal origin,* and *epistemic context*. The next section introduces a formal language to do just that, without overthrowing the valuable parts of the classical understanding.

Language and Semantics

We formulate our theory in a **many-sorted first-order language** with equality. "Many-sorted" means that our domain of discourse is divided into several disjoint sorts of entities, each with its own identity predicate (we may still use "=" symbol, but it is always sort-specific, never equating entities of different sorts). This allows us to distinguish, for example, **individuals** (ordinary concrete objects), **times**, **locations**, **events**, **tags**, and perhaps **kinds or natures** as separate sorts. By having multiple sorts, we can avoid category mistakes – e.g. asking if a time = a person is simply ill-typed and not formulable in the language. Each sort comes equipped with its own equality relation, which is assumed to satisfy the usual axioms of identity (reflexivity, symmetry, transitivity, and Leibniz's Law restricted to that sort). In particular, if "=" with a subscript denotes equality in a given sort, then for any sort σ and any x, y of sort σ, we have: $x =_\sigma x$ (always), and if $x =_\sigma y$, then for any well-formed formula φ in that sort, φ(x) implies φ(y) (the substitutivity schema). These are just the ordinary first-order logic with equality axioms, but **sorted** [17]. This sorted approach is conservative over standard logic – it doesn't change any inferences about a single sort's objects, it just prevents cross-sort confusion.

Within this language, we introduce three key primitives that extend the resources for talking about identity:

1. **Provenance Tag Function:** We include a function symbol (or a function *from* individuals to the Tag sort) which we might write as **Tag(·)** or just τ(x). For each individual object *x* (of sort Individual), τ(x) yields a special entity of sort **Tag** (or "origin-tag"). Intuitively, this tag represents *the provenance or origin* of x – something like the event or process by which x came into existence, or a unique marker of that coming-to-be. The fundamental semantic postulate is: **two individuals are numerically identical (the very same entity) if and only if they share the same tag**. In other words, for any individuals x and y, we do not say x = y in the usual sense unless τ(x) = τ(y) (equality in the Tag sort). If τ(x) = τ(y), we will also say x and y have the same "origin" or belong to the same genesis. The theory will ensure that τ(x) = τ(y) is an equivalence relation mirroring object-identity on the individual sort.

In fact, we can take *identity of individuals* as a *defined notion*: "x is-identical-to y" can be defined to mean τ(x) = τ(y). This provides a *truth-maker* for numerical identity: it's not just a primitive fact that x = y, but rather there is an underlying tag that both have by virtue of how they came to be. A simple way to think of tags is as **surrogate keys** or GUIDs (globally unique identifiers) assigned at creation – widely used in databases and information systems to track whether two records refer to the same entity. However, in a philosophical sense, we intend τ(x) to capture what philosophers like Saul Kripke and others call the **origin essential** property of an object (Kripke 1980). For example, an organism's tag might encode its conception event (parents etc.), an artifact's tag might encode its initial production. If one attempted to create a "duplicate" of an object ab initio (like a perfect copy emerging independently), in our theory that duplicate simply would have a different tag, and hence would not be identical to the original despite any qualitative similarity. We thereby honor the intuition that *origin matters* for identity (Kripke's dictum that a person born of different parents would not be the same person, etc.), but we do so by building origin into the formal identity criterion. We stress that **tags are not haecceities in the spooky metaphysical sense** of non-qualitative thisness properties; rather, they are logical placeholders for actual origin events or chains of custody. They are "logical atoms" in the Tag sort that ground identity without themselves having an internal structure in the theory (though informally we understand them via origins). This approach is similar to hintikka's world-lines or to Quine's idea of reinterpreting predicates to treat an equivalence class as an individual [18], but here we explicitly include the tags in the ontology for clarity.

2. **Continuity Predicate:** Next, to handle persistence through time, we introduce a binary predicate *C*(x, t) which says that individual *x* **exists/occurs at time t** (we could also think of it as x being "present at" t). Using this, we can define what it means for an individual at one time to be "the same continuant as" an individual at another time. One way is to introduce a tertiary relation *Cont*(x, t1, y, t2) meaning "x at t1 is continuous with y at t2". However, we can simplify by treating an individual's presence at time as a state, and say that if x and y are the same continuant, any time associated with one corresponds to the other. In practice, it is easier to introduce a notion of **world-line** or history for each individual: think of it as an equivalence class of (object, time) pairs. Formally, we can define an equivalence relation on pairs (x, t) and (y, t′) that holds iff x and y are stages of a single continuing entity. The key new ingredient is an **anti-branching policy**: in scenarios where one "continuant" could branch into two, a policy will decide either that one branch gets the original identity or that the original terminates and new ones begin. We do not bake one specific policy into the logic; instead, the logic can accommodate a parameter or context specifying the policy. For example, a policy might be "whichever branch has more of the original's parts/character" (a closest continuer rule), or it might be context-dependent (legal convention might say the original company is the one retaining the original name in a spin-off). The important thing is that in any *given model*, once a policy is fixed, **continuity yields an equivalence relation** partitioning the set of object-stages across time. If object *A* at time 1 and object *B* at time 2 fall in the same equivalence class, we say informally that B is the **continuation** of A (or A persists as B). This relation of being on the same world-line will be reflexive, symmetric, and transitive (once no-branching is assumed, any stage is "continuous with" itself, and if A at t continues as B at t′ and B at t′ continues as C at t″, then A at t continues as C at t″). Persistence through time, under this approach, is not simply identity (since, for example, the person-stage at age 10 and the person-stage at age 40 are not literally the same *stage* or same momentary object), but they are linked by continuity. We do not say an individual *changes* its identity; rather, we say that statements like "X at t = X at t′" are shorthand for "there is one continuing individual (one world-line) that has X at t and X at t′ as stages." In our formal theory, we might not need a separate symbol for

the continuity equivalence if we instead introduce a sort for "histories" or continuants. Alternatively, we can introduce a predicate **SameLine(x, y)** meaning x and y (which may be stage slices) belong to one continuous history. The exact formalization can vary: one convenient method is to introduce a function that assigns to each individual a set of time-indexed stages or something akin to a "history ID." For simplicity, let us assume *C*(x, t) is primitive (read "x exists at time t"), and we add axioms to enforce that for any x, the set of pairs (x, t) is linearly ordered by time (no branching) unless branching is explicitly allowed as indeterminate. If branching is allowed, then *C*(x, t) could hold for multiple x with overlap, but then identity claims will be marked indeterminate. We will formalize this in the axiom section. Semantically, one can think of each individual having a "timeline" and continuity means sharing that timeline. This addresses the problem of diachronic identity by grounding it in actual continuous existence rather than mysterious endurance without criteria. The **anti-branching policy** is made explicit: rather than pretending identity magically knows what to do when one becomes two, we say upfront that one must stipulate a criterion (or accept indeterminacy). This approach draws on ideas from the literature: for example, Nozick's closest continuer (1981) and Lewis's insistence that identity *does not branch* (which he addressed by perdurantism – branching corresponds to overlapping person-stages, not one person becoming two) [19] . Here we simulate the effect by logical rules: if branching occurs, either arbitrarily pick one branch as continuation (policy) or deem the identity facts as not fully determinate until policy resolves it.

3. **Constitution and Aspect Operator:** Finally, we introduce an operator to handle cases of "sameness in a certain respect" – the paradigm being the statue and the clay, or generally an object and the material that constitutes it, or an object and a role it plays. We denote this as something like **X ≡_K Y** or use a binary predicate **Const(X, Y, K)** meaning "X and Y are the same **K**" (same *in kind* K, or X constitutes Y as an instance of kind K). Another approach is a unary operator that can modify a term with an aspect, often called a "qua" operator in literature (Fine 1982 introduced the notion of *qua-objects* to deal with these issues). For example, we might say **a = b (qua F)** to mean "a and b are the same F." In our formal language, one way is to allow terms of the form *x:K* (x-as-a-K) and introduce an equality on those, or a predicate that links x and y under K. To keep things straightforward, let's use a ternary relation Const(x, y, K) where K could be a sort or a tag indicating a category or aspect. If K is a sort (like "Person" or "MaterialObject"), this relation might assert that x (of sort Person) and y (of sort MaterialObject) coincide in that y constitutes x (the person) at the moment. Or if K is more like a context label (like "from the perspective of material constitution"), it indicates an aspectual equivalence. The *semantics* of this constitution/aspect relation is that it is **reflexive** (everything is trivially the same as itself under an appropriate aspect) and typically **symmetric** (if x constitutes y, then y constitutes x, in the sense that they share all relevant features in that context, though not necessarily share identity). But constitution is **not** an unqualified equivalence relation across the board: it is not transitive in general, it is not a replacement for identity. Rather, it obeys a restricted form of substitutivity: one can substitute x for y in formulas that are within the same aspect or context K, but not outside of it. For instance, if we know Statue ≡_Material Clay (the statue and clay are the same material object), then any predicate about mass, location, shape (physical properties) might hold equally of both at that time. But a predicate like "is created in 2021" might apply to the statue and not the clay (maybe the clay existed earlier), which is fine because that's not purely within the *Material object* aspect (it involves historical properties). The logic will enforce that from "x ≡_K y" one cannot infer "x = y" (numerical identity) or mix predicates from outside the K-context without checking conditions (this is our **NoMix** rule informally). One way to implement this formally is to treat "=" as always sort-specific and to treat x ≡_K y as saying some projection or feature of x and y is identical. For example, if K = "material", we could say the **matter of** x = the **matter of** y. In statue and

clay, the statue's matter = the clay's matter, which is true; but the statue and clay themselves are not identical because they have different form, different persistence conditions, etc. Our constitution operator generalizes this idea: it licenses us to say when two things can be regarded "as the same in certain respects" without ever conflating their identities. Another important use-case is in **opaque contexts**: we might treat, say, the belief of Lois Lane about "Superman" and "Clark Kent" as involving an aspect (the mode of presentation or guise under which the man is considered). If Lois knows Clark as "Clark" but not as "Superman," we could represent that by saying Lois's belief context is sensitive to the aspect (name), and **Clark ≡_{cognitive} Superman** may fail even though Clark = Superman in reality. Or better, we do not even put Clark and Superman (the person) as distinct individuals; rather, we treat "Clark" and "Superman" as two aspects or guises of the same individual. Then substitutivity fails because Lois's belief is keyed to one guise. This however gets into a more complex territory of structured propositions; for now we note that the **NoMix** constraint in our theory will forbid naive substitution of one co-referential term for another across such aspectual contexts [3] .

To make the semantics more precise in plain prose: the many-sorted domain might include:

- **Individuals (I)**: things like persons, ships, statues, lumps, etc. Each individual x has a tag τ(x) which is an element of sort **Tag (T)**. Two individuals are the *same* individual iff they have the same Tag. In any one model, we can think of each Tag as corresponding to a unique individual's identity – effectively, tags partition the individuals into identity-equivalence classes. (We could even choose to model individuals as ordered pairs (tag, state) if we wanted, but that's an implementation detail.)

- **Times (Time)**: moments or intervals. There is a temporal ordering and maybe metric, but logically we just need a precedence relation or something linear.

- **Events (E)**: possibly if we want to model origins as events, each tag could be associated with an event of coming-into-being in E. For simplicity, one can collapse Tag and Event sorts by saying each origin event yields a unique tag, and vice versa. But keeping Tag as an abstract key is fine.

- **Kind or Aspect identifiers (K)**: This could be a sort that labels different aspectual contexts or type of constitution relation (like "material", "formal", "functional", "divine nature", etc). Alternatively, K could be a finite set of relations built-in, not necessarily a domain.

Now, **vagueness and indeterminacy**: We anticipate that sometimes the provenance of an object or the boundaries of continuity might not be sharp. For example, an object might form gradually (when exactly does a heap become a heap, or when exactly did this river begin? origins can be vague). Likewise, continuity might be vague (if a transport device gradually disintegrates and duplicates, is there a fact of the matter which is the original?). To handle this, we adopt a **supervaluationist** semantics for identity statements: we consider a range of admissible precisifications (precise ways to assign tags and continuity relations to borderline cases), and we say a statement is true if it is true under all precisifications, false if false under all, and otherwise neither (which can be treated as indeterminate). Crucially, this does *not* make our logic non-classical – we can stick to classical proof rules – because we are not treating "x = y" as itself a vague predicate with a truth-value gap in the logic. Instead, we treat the vagueness as *metalinguistic or semantic indecision*. In practical terms, we will not delve deeply into this in the axiomatics, but we assume that for any situation where it is indeterminate whether x and y share a tag or share a continuity line, the correct approach is to evaluate identity claims under all sharp assignments and see if the result comes out

uniformly. This approach, following Fine (1975), **allows us to preserve classical inference rules** (no violations of Leibniz's Law or contradictions) while acknowledging that sometimes we cannot definitely say whether x = y [8] [20]. The logic, in other words, remains two-valued and classical at the level of completed models; vagueness is handled by quantifying over multiple classical models. This is standard supervaluation: it vindicates the law of excluded middle (either x = y or x ≠ y is true in the supervaluational sense, even if neither disjunct has a determinate truth on its own – well, actually in supervaluation, if neither has determinate truth, the disjunction might still be true by being true in all precisifications that one or the other holds, but one has to be careful with higher-order vagueness). For our purposes, we simply note: the proof system can remain classical because any proof will be under the assumption of some precise model. This means we do not countenance "vague identity" as a fundamental notion (we side with Evans 1978 that treating identity itself as internally vague leads to trouble), but we allow that it can be *indeterminate which identity statement is true* due to underlying vagueness in inputs (origin, continuity) [21] [22]. This is much like saying a sorites series leaves it indeterminate which grain's removal stops it being a heap, but we don't say "heapness" is a truth-value-gap predicate in the logic, we just say the term "heap" has borderline cases.

To summarize the language and semantics:

- It is many-sorted: key sorts include Individuals, Tags, Times, (possibly Events, Locations, Kinds).
- Each sort has its own identity = which is an equivalence relation satisfying the usual laws within that sort.
- There is a function τ: Individual → Tag (for each individual get its origin tag).
- There is a predicate C(x, t) meaning "Individual x exists at time t" (for persistence).
- There is a relation or operator for constitution/constitution-as: we will treat it as an infix "≡" annotated by a category, or as Const_K(x, y).
- Additional structure: ordering on Times, maybe a function mapping Tag to an Event (the origin event), etc., could be added but not essential.

With these in place, we now turn to the axioms that formally govern these primitives and ensure they fulfill the intended role.

Axiomatics

We present the axioms of the theory in plain language (each could be formalized as a sentence or schema in the many-sorted language). The axioms are grouped by the three explanatory bases: provenance, continuity, and constitution (plus a constraint about mixing levels). Throughout, we take it as given that ordinary first-order logic axioms for equality hold within each sort.

Axioms for Provenance (Origin-Tag Identity):

1. **Tag Equivalence:** For any individuals x and y, if τ(x) = τ(y) (their tags are equal in the Tag sort), we say x and y are **numerically identical**. This is effectively a definition: "Id(x,y) := [τ(x) = τ(y)]." We then require that this identity behaves as an equivalence relation on Individuals. Since equality in the Tag sort is already an equivalence relation (reflexive, symmetric, transitive) by the built-in logic of equality, it follows immediately that Id(x,x) holds, Id(x,y) implies Id(y,x), and if Id(x,y) and Id(y,z) then Id(x,z). In other words, **sharing the same tag** is an equivalence relation on individuals. We correspondingly impose **Leibniz's Law for Id**: if τ(x) = τ(y), then for any property or formula φ in the

language of individuals (that doesn't mention higher-order things like tags explicitly), φ(x) implies φ(y). In practice, we might simply enforce that the identity predicate in the Individual sort is interpreted as Tag equality. This ensures substitutivity: within the Individual sort, identicals (same tag) are indiscernible [1]. Note that because our language allows talking about tags, we must be careful: we do *not* allow substituting individuals across an expression like τ(x) = t unless we know those individuals share the tag. But that's just normal substitutivity anyway. So in summary, these axioms make provenance identity coincide with classical identity on individuals.

2. **Uniqueness of Tags:** Each individual has exactly one tag (the τ function is total), and if τ(x) = τ(y), then x and y are *the same individual* (so τ is many-to-one mapping collapsing exactly those individuals that are identical). We can formalize: ∀x∀y [τ(x) = τ(y) → x =_I y]. This essentially says tags are unique identifiers – no two distinct individuals share a tag. Combined with the definition of Id above, it becomes a tautology. But conceptually, it means tags carry all information needed to tell individuals apart. This axiom is what gives the system the "no haecceitism problem" reassurance: a tag is not an arbitrary label that two different things *could* share; if they did share it, they wouldn't be two different things but one. In model-theoretic terms, τ is an injective function from the Individual sort into the Tag sort (perhaps even bijective if every tag actually comes from some individual).

3. **Tag Indiscernibility** (optional specific form): If we wanted, we could also add an axiom schema that any primitive or defined property of individuals that is *preserved under* $=_I$ must actually be a function of the tag. But this might be too strong or unnecessary. Intuitively, however, if x and y have the same tag, anything that does not explicitly involve time or constitution (i.e. anything about their identity as objects) will be the same for x and y. We likely don't need this explicitly since substitutivity covers it.

Thus, provenance-based identity fully respects the **formal requirements**: identity by tags is reflexive (trivially, since each x shares its tag with itself), symmetric, and transitive, and obeys substitutivity for well-typed formulas [1]. We thereby satisfy the classical meta-constraints in the domain of individuals. The difference is that now we have a substantive criterion: *x = y if and only if x and y originate from the same source*. In scenarios where an object is "duplicated" (two distinct objects with identical qualities), those two will simply have distinct tags, so we correctly say they are not identical. This handles puzzle cases like duplication or imaginative sci-fi cloning: no matter how indistinguishable the two are, they are numerically distinct unless one literally *turned into* the other (which would entail same tag). If a duplication machine creates B as a copy of A, then B gets a fresh tag distinct from A's tag, and hence A ≠ B. In some metaphysical debates this is sometimes questioned (could two distinct origin events yield the same individual? No, by definition, since identity is one per origin). Our theory essentially builds in a form of **origin essentialism**: an individual cannot "change" its origin or have come from a different origin – if it had, it would be a different individual (different tag) (Kripke 1980). In the logic, this can be expressed by the necessity of identity: necessarily, if τ(x) ≠ τ(y) then x ≠ y. But we won't go into modal aspects here.

Axioms for Continuity (Persistence):

1. **Continuity Equivalence (under policy):** We fix a criterion that in any model, determines an equivalence relation on (Individual,Time) pairs, intended to represent the *same continuing entity*. Formally, we could introduce a predicate ~ such that (x,t) ~ (y,t′) means "x at t and y at t′ are on the same world-line." The axioms would be: (a) ~ is an equivalence relation on the set of all (x,t) for which C(x,t) holds (reflexive: any (x,t) is ~ itself; symmetric and transitive accordingly). (b) **Anti-branching**: if (x,t) ~ (y,t′) and (x,t) ~ (z,t′) (the same stage x at an earlier time continues as two different individuals y

and z at a later time t′), then we must have y =_I z – i.e. those later individuals are in fact the same individual. In other words, one history cannot fork into two distinct individuals. This enforces the no-branching condition. However, we allow that some situations might violate this as a matter of empirical description (like identical twins splitting from an embryo). In such cases, rather than violate the axiom, we will say our *description* must decide which way we count identity. If no decision, we treat one or more identity statements as indeterminate. In a supervaluational view, the different choices of which twin gets the "original" tag are different precisifications. The formal axiom might thus be conditional: If (x,t) ~ (y,t′) and (x,t) ~ (z,t′) *and* (some policy condition holds at t′), then y = z. The "policy condition" might be something like: if two candidates exist, either pick one (then the other continuity relation doesn't hold), or if none can be chosen, then ~ doesn't relate the earlier stage to any later ones (meaning the original ended). We can simulate this by saying whenever continuity would not be functional, the statement "there exists a y such that (x,t0) ~ (y,t1)" is neither true nor false until disambiguated. Rather than complicate axioms with that, we assume a given model comes with a branching-resolved world-line relation. Essentially, we ensure each equivalence class under ~ looks like a "tree" or rather a line: at most one member at any later time.

(c) **Temporal coherence:** if x and y are on the same world-line and both exist at some time t (like x = x's stage at t and y's stage at t coincident), then x and y must actually be the same individual. Typically, in a no-branching scenario, an equivalence class of ~(continuity) will have at most one individual at any given time. We enforce that: ∀x,y,t [C(x,t) & C(y,t) & (∃(x0,t0) such that (x0,t0) ~ (x,t) and also ~ (y,t)) → x =_I y]. In simpler terms: one continuant cannot have two different individuals existing at the same time – otherwise it's two continuants.

To ground these axioms intuitively: we are saying each individual has a *story through time* (if it persists at all), and that story cannot branch into two individuals nor can two stories fuse into one unless the objects fuse (which would be represented as one ending and being taken over by another's tag perhaps). If a fusion occurs (two things merging), under our approach, likely one or both original tags cease (the original individuals die) and a new tag is born for the fused entity, or one tag prevails if we consider one absorbed into another. The policy might dictate that.

One could formalize an "only one survivor" rule for fission and "only one antecedent" for fusion. But a flexible way is just to rely on tag assignment consistent with observed continuity.

Thus, continuity plus anti-branching yields a well-defined **persistence equivalence**. This equivalence is what underwrites ordinary talk of an object being "the same" at different times. In our system, strictly speaking, if one asks "Is object A at time t the same as object B at time t′?", the answer is yes if τ(A) = τ(B) (same tag) *and* there is a continuity chain connecting A's state at t to B's state at t′ without violation of policy. If τ matches but continuity fails (e.g. someone died and then a qualitatively similar person appears later with the same tag by some weird scenario), that scenario cannot happen if tags truly represent origin (the later identical person would have a different tag). If continuity holds but tags differ, that means we have distinct individuals that are continuous in some way (which usually can't happen unless perhaps two individuals merge – but then one tag would presumably drop out). So normally, tag equality and continuity will coincide for well-behaved cases.

However, in duplication thought experiments, continuity might suggest two possible futures for one past person, but tag identity can only follow one. Our system's resolution: both future copies have continuity from the original in a qualitative or causal sense, but our **semantic policy** will not allow both to count as

numerically the same; it either designates one as the successor (and gives it the original tag, the other gets a new tag) or considers the original's identity to have ended (the original tag retires) and both copies get new tags. Which it is can depend on context. The important point is, no formal contradiction arises – we do not say "A = B and A = C but B ≠ C" because that scenario is ruled out by anti-branching or by avoiding giving A's tag to both B and C [23]. Instead, the situation of uncertainty is represented by potentially multiple admissible assignments of tags (one assigning the original tag to B, another to C, another giving to neither). Under supervaluation, a statement like "B is the same person as A" might come out indeterminate if in some precisifications B inherits the tag of A and in others C does. But logically, in any single precisification (model), the statements are classical (either true or false with a single assignment of tags and continuity).

To capture **indeterminacy** formally, one could include a predicate for "Indeterminate(x = y)" or just handle it semantically. The latter is fine: we say if an identity question can't be answered due to branching, the theory does not assert either x = y or x ≠ y as a theorem, leaving it open.

Thus, our continuity axioms allow us to talk about persistence *without* forcing arbitrary verdicts in problematic cases. We do not have to declare, as a law, that "if psychological continuity then identity" or any specific criterion; we leave it abstract but structured by whichever policy is chosen.

Axioms for Constitution/Aspect:

1. **Constitution Reflexivity:** For any object x and appropriate aspect K, x ≡_K x (an object is trivially the same as itself under any aspect it trivially exemplifies). For example, a clay lump is the same material as itself. This axiom is just a sanity check: constitution is at least reflexive (some would argue it's not always reflexive if the aspect isn't applicable – but we assume each object has itself as a trivial constitution).

2. **Constitution Symmetry (for given aspect):** If x ≡_K y, then y ≡_K x. This usually holds because if x constitutes y as a K, then y is constituted by x as a K, implying a mutual dependence in that aspect. For instance, if a statue is constituted (materially) by a lump, the lump constitutes the statue in that same material sense at that time. We will assume symmetry to simplify (some accounts consider constitution a one-way dependence, e.g. lump constitutes statue but statue does not constitute lump, because lump can exist without statue; however, we are formalizing a relation of sameness-in-K, which should be symmetric equivalence in that context even if the **modal dependence** is asymmetric. So yes, *sharing all K-properties* is symmetric).

3. **No Cross-Context Substitution (NoMix Constraint):** If x ≡_K y, then for any formula φ(z) that is not purely about aspect K (i.e. that has predicates outside those valid in context K), one cannot infer φ(x) ↔ φ(y). Substitution is only licensed for formulas describing the K-aspect. Concretely, if K = "material makeup", then from Clay ≡_material Statue, we can infer they have the same weight, same shape (assuming shape is purely material), etc., but we cannot infer they have the same temporal duration or same artistic value, etc. This "NoMix" rule is not a single axiom but a restriction on how the inference rules apply. We can implement it in axiom form by introducing sortedness or typing for aspects: say we have special predicates that only apply to things considered as K. One formal approach: treat ≡_K as a sort of identity *within a typed domain K*. We then say it obeys Leibniz's Law for all predicates of that *aspect sort*, but not for others [17]. For example, if F is a predicate like "...is made of clay" (a material property), and Statue ≡_material Lump, then F(Statue) ↔ F(Lump). But if G is "...was created in 2022" (a historical property), we do not assume G(Statue) implies G(Lump),

because that crosses out of pure material talk. In practice, we won't enumerate such axioms here, but they are understood: constitution-same objects share all **categorical properties relevant to that aspect** but may differ in others. As a specific **NoMix** principle, we forbid any inference of the form "x ≡_K y and x has property P → y has property P" unless P is explicitly marked as a K-level property. This blocks the classic mistake of inferring identity: e.g., "the statue is the same material as the clay, and the clay existed since yesterday, so the statue existed since yesterday" – that inference is disallowed because "existing since yesterday" is not a material property but an identity/time property.

4. **Non-Transitivity across distinct bearers:** Constitution should not generally be transitive, especially across time or individuals. For instance, if A ≡_K B (A and B are the same K-wise) and B ≡_K C, it does not necessarily follow that A ≡_K C, if A, B, C are distinct entities. A concrete example: a bronze sphere is constituted by a lump of bronze (A ≡_material B). That lump of bronze is later reshaped into a bronze cube (C), and at that later time the same lump constitutes the cube (B ≡_material C at that time). The sphere and the cube are not the same object in any sense; they existed at different times, and one did not turn into the other directly (though they shared matter). Transitivity would wrongly conclude sphere ≡_material cube, which is not true – they never coexisted; we only talk about constitution at a given time. So to be precise, constitution ≡_K is typically defined at a time or in a context. We should index it by time if needed: e.g., at time t, statue ≡_material lump. At a later time t′, maybe not if the statue is destroyed. So if needed, we assume constitution statements are implicitly temporally qualified (or eternal if both exist throughout an overlap).

Given that, we ensure no chain of constitution across distinct things yields an identification. Each ≡ is local to a context/time. We might add an axiom: If x ≠_I z (not literally same tag) and x ≡_K y and y ≡_K z, then typically either y is an overlapping part or time. Actually, let's keep it simple: we do not allow repeated application of ≡ to simulate =. This could be a meta-rule: do not conclude x = z from x ≡_K y and y ≡_K z for any K (even if same K or different). Only the trivial case where all three are actually one individual (like x ≡_K y and x ≡_K z could imply y ≡_K z if all at same time context, but that is fine because then likely all three share matter at that moment).

1. **Kind Predication Constraint:** This is a specific NoMix rule: if a is of kind F and b is of kind G and a ≡_K b, one cannot infer F = G or that a is of kind G, etc. For example, from "Clark Kent ≡_person Superman" (which is trivial since they are the same person – maybe a bad example, better: "the statue ≡_material the lump (which is a lump of clay)", one should not infer "the statue is a lump of clay" as a literal identity of kinds. The statue *is made of* clay, but it is not a clay lump by kind. So we keep track that being the same in one aspect does not collapse the categories. This prevents arguments like: The clay is a lump (a mass of clay), the statue is the same material as the clay, therefore the statue is a lump of clay – which is usually considered a category mistake. Our logic won't allow concluding an Individual's sort from the aspect identity alone.

In effect, the **NoMix** and related constraints preserve the distinction between **numerical identity** and **constitution**. We have two relations in the system and we never let one convert into the other improperly. This addresses an objection sometimes raised to constitution views: that saying "the statue is the same as the clay in some sense" is just relative identity by another name (Geach 1967). Our formal constraint ensures it doesn't behave like normal identity across contexts – it cannot violate classical logic because we never allow an inference that would produce a contradiction with classical identity. The statue and clay

remain distinct individuals with different tags (so all classical identity laws apply to them: they are not equal, period), but we simply have an additional relation linking them, which is carefully constrained.

Finally, we might add a technical axiom about **existence conditions**: If x ≡_K y, usually x and y must co-exist (at least during the interval of constitution). You cannot have constitution without the two being around at the same time or one being a part of the other's existence. For instance, the statue and clay exist simultaneously while the statue exists. If statue is destroyed, it no longer is constituted by the clay (which now perhaps exists as a shapeless lump). We can formalize: x ≡_K y implies there is some time t such that C(x,t) & C(y,t) (they coincide in time). This ensures we don't get weird across-time constitution claims (which anyway would be more like continuity).

Thus the axioms for constitution/essence talk can be summed up: constitution is a **restricted equivalence** that is sort-of an equivalence relation on things considered under a certain aspect, but it is explicitly *non-compositional* across different aspects or outside that aspect's domain [17]. It never allows one to prove identities or equalities outside its scope. We might say it's an equivalence in a substructure of the model (like within a time slice and category).

No-Mixing (Overall Principle): To encapsulate: No formula that mixes the provenance identity and constitution relations inappropriately is allowed to shorten the gap. For example, one cannot use a chain of continuity and constitution to jump to a conclusion of identity. This can be thought of as a kind of type safety: identity is a notion for individuals; constitution is a notion relating perhaps individuals of different categories; continuity relates individuals over time. There is no general law bridging these except trivial ones (if x = y then certainly x ≡_K y for any K trivially, and if x = y then x is continuous with y always). But not vice versa.

To illustrate these axioms at work: - In the Ship of Theseus puzzle, we have one original tag for Theseus's ship. As the planks are replaced, continuity is maintained along that ship, and anti-branching says we can't have two ships sharing the original tag. So what happens when the planks are reassembled into another ship? In our theory, that reassembled ship must get a **new tag** (fresh provenance) because its coming-into-being event is the assembly later. Therefore, it is *not* identical to the original ship (different tag), even though it is made of the same planks. However, we can say the reassembled ship is constituted by the same material as the original at an earlier time, or shares the same structure, etc., but not the same identity. Thus only the continuously repaired ship keeps the original tag and thus is identical to Theseus's ship [5]. This matches one common solution and avoids transitivity issues entirely, since we never had original = reassembled in the first place. The laws enforced that.

- In a fission case (like a person splitting into two), anti-branching ensures not more than one future individual can carry the original tag. If a policy is in place to pick one (like perhaps arbitrary or contextually "the left branch is original"), then that one gets the tag, the other gets a new tag, and so identity claims follow that assignment. If no policy, our semantics would say the statement "FuturePerson1 = Original" is indeterminate and likewise for FuturePerson2, but we will never derive a contradiction because we won't assert both or anything.

- In a statue/clay case, tag of clay and tag of statue are different (statue's tag originates at sculpting time from the clay but it's a new entity in terms of form). But when co-located, Statue ≡_material Clay holds. The axioms let them share physical properties (mass, spatial location) because those are aspect-specific, but do not let "is clay" or "began existence in prehistoric times" transfer from clay to

statue. The statue's tag has a birth when sculpted, so the statue doesn't exist before then, satisfying intuition while still letting us say in a loose way "the clay and statue are the same thing *in the material sense*" without committing identity. This resolves the puzzle without confusing our logical identity predicate.

In formal verifications, one would include all the above axioms as either axioms or theorems from definitions. The system can then prove various facts, e.g. if an object's matter is replaced gradually, it can still have the same tag, etc., but if it's duplicated, one copy must differ in tag.

One last axiom worth stating is about the relationship between these notions and standard first-order logic:

- **Conservativity Axiom:** There are no bridge axioms that identify or collapse sorts. For example, we do **not** have any axiom saying something like "Every tag is a unique property of some individual" or "for each tag there is at most one individual" beyond what's needed for identity. We also do not allow tags to have rich internal structure that could re-introduce paradox (they're atomic). The result is that if you ignore all the new symbols (τ, C, ≡, etc.), the remaining theory on just individuals with = is just classical first-order logic with no extra constraints (thus conservative). If you consider formulas that do not mention τ, C, ≡, their truth is unaffected by adding our axioms (hence any model of a pure identity theory can be expanded to a model of our theory by suitably adding tags and such). This means our theory does not secretly impose any unusual condition on standard domains; it only adds new distinctions.

Summarizing: provenance identity (same tag) gives us a robust equivalence relation capturing numerical sameness; continuity (with no-branching) gives us a way to relate stages and talk of persistence in a way that, when branching is not resolved, doesn't lead to contradiction but to indeterminacy; and constitution (aspectual equivalence) lets us handle composite objects and intensional identity inferences by carefully controlling where substitution is allowed. We have essentially built an **identity calculus** that distinguishes different "sameness" relations: = *for absolute identity,* ≡ *for aspectual sameness, and continuity for persisting existence.* Each obeys laws akin to identity (equivalence) in its own domain, but importantly, they are separate and never conflated.

Meta-theory

Having laid out the theory's syntax and axioms, we now address some meta-theoretical properties. First, we claim that our theory is a **conservative extension** of standard first-order logic with equality (FOL=) on the ordinary domain of individuals. By conservative extension, we mean that any statement in the language of ordinary first-order logic (with just the usual "=" on individuals and no mention of tags, continuity, etc.) that is provable in our theory was already provable in pure logic or follows from the axioms of equality alone. Equivalently, if you restrict attention to sentences that do not involve the new primitives (τ, C, ≡, etc.), our theory does not produce any new theorems about them beyond what classical logic would. This is ensured by the fact that our additional axioms are all sort-confined and do not identify or equate distinct sorts. For example, we did not assert anything like "for every individual there is a unique tag" as a separate principle – we could, but that just assures existence of tags, which doesn't affect pure individual-only formulas except by consistency strength.

Moreover, the theory is **internally consistent** (at least as consistent as classical logic). We can provide models to demonstrate consistency. One simple model (Model 1) is a universe where: the Individual sort is,

say, a set of people {Alice, Bob, ...}, the Tag sort is maybe identical to the set of individuals (each individual is its own tag; τ is identity mapping). Time could be a single point if we don't care about change, or a small set, and C(x,t) holds for all x at t (everyone exists at that one time). Constitution can be empty or trivial (no two different individuals are constitution-related except itself). This trivial model satisfies all axioms: each person's tag is itself, so provenance identity just says x = y iff x = y (no surprise); continuity is trivial as there's only one time or no interesting branching; constitution is trivial. So consistency at least is as good as classical first order (since we basically interpret new symbols in a degenerate way that doesn't cause contradiction).

More interestingly, we can construct a **model with duplication** (Model 2) to illustrate how the system handles it without contradiction. Imagine a scenario with times t0 and t1. At t0 we have one individual A (perhaps representing an amoeba before splitting). At t1, we have two individuals, B and C, which are the "offspring." We want to model that A fissioned into B and C. In our theory, we cannot have A's tag given to two individuals. So how do we represent this? Either (a) we say the model's assignment of tags chooses B as having the same tag as A (thus B is the "continuation" of A), and gives C a new tag (so C is a new individual that came into existence at t1 independently, even if physically it was part of A). This satisfies anti-branching: A's tag = B's tag, and A's tag ≠ C's tag, so no one tag is on two distinct individuals. The continuity relation can reflect that A at t0 is continuous with B at t1 (same world-line), but not with C (C is a new line). So in this model, the theory will conclude A = B (they share tag, just at different times, effectively the same persisting entity) and A ≠ C. This matches a stipulative choice that B is the "same organism" as A and C is a twin or clone. Or (b) we could say neither B nor C carries A's original tag (i.e., perhaps A's tag "expired" with A and two new tags were minted for B and C). Then A simply has no continuation; B and C are both new. In that model, the theory would say A is not identical to either B or C (different tags), and just treat it as death and birth of two new ones. That's a different outcome (and indeed some philosophers say in fission the original ceases, two new ones form). Both models are allowed by our axioms (depending on the policy). What's important is: in neither case do we derive a logical contradiction (like B = C or A = B = C, etc.), and yet we can still talk meaningfully about continuity: e.g. in model (b) we might still have a continuity relation linking A's stages with B's and A's with C's, but because we allow indeterminacy, maybe we interpret continuity loosely. However, if our logic demands an equivalence, we can't call it an equivalence if A is connected to both B and C. So in a single model, we likely implement the anti-branching by just not connecting A to both in the same model. The scenario (b) would mean A's continuity didn't extend past t0 (it ended), and two new lines begin at t1. Each of B and C have their own tags.

Thus, in any such model, all axioms hold, and we simply have different **possible outcomes** for identity claims in a fission scenario, none of which break the formal rules. This shows the theory is consistent with the phenomenon of duplication: it doesn't explode. In contrast, naive identity theory would have trouble describing fission without either arbitrarily disallowing it or ending up with indeterminacy that it cannot formalize. Our system formalizes it via the policy or tag assignment approach.

We should also mention **completeness and compactness**: Since our theory is essentially first-order (many-sorted first-order, which is equivalent to single-sorted first-order with unary predicates partitioning sorts), and since we haven't introduced any higher-order or modal operators, the usual meta-theorems apply. In particular, if the axioms we gave are all first-order sentences, then by the completeness theorem, if something is semantically entailed it is provable. We don't expect any exotic behavior like needing second-order logic. The presence of an anti-branching condition might raise eyebrows – because anti-branching can be seen as a second-order condition (it's like a uniqueness of extension condition). But we have formulated

it in first-order terms (∀(x,t),(y1,t′),(y2,t′) if (x,t) connects to both (y1,t′) and (y2,t′) then y1 = y2, which is first-order expressible if you treat times and individuals as objects). So it's fine.

Thus the theory is axiomatically complete relative to first-order logic with these symbols. It is also likely **decidable only in fragments**. The whole theory is probably not decidable, because even first-order logic with one equivalence relation isn't trivially decidable (and adding multiple sorts and functions likely yields an undecidable theory in general). However, certain fragments might be decidable. For instance, if we restrict to a fixed finite number of time points or to a monadic fragment, etc., one could get decidability. But in full generality, since this theory can encode a fair amount of structure (especially if we consider arbitrarily many objects and times, etc.), it will likely be undecidable, like typical first-order logic with equivalence relations (which can simulate arithmetic in some cases).

Nonetheless, for practical use, often we are interested in finite scenarios or ones amenable to algorithmic checking (like verifying identities in a database with no cycles).

We can provide a simple **model existence result**: Given any set of "histories" you want, you can construct a model. For example, suppose we want a model with two individuals that split into four by time t2. We can do it by introducing appropriate tags and not connecting them illegally. In general, one can always model a scenario of branching by either picking a survivor or terminating identity. So consistency is easy.

From a more abstract perspective, our theory can be seen as a combination of several disjoint equivalence relations (on tags, on stages, on aspect pairs). This hints that the theory might even be broken into components that are separately complete and then combined by union of theories (not interfering with each other's languages except via a few bridge constraints). The conservativity ensures we didn't break completeness of standard logic: any purely individual-level inference is still governed by plain logic.

One important meta-theoretical aspect is how our theory handles **indeterminacy** logically. Strictly speaking, classical first-order logic doesn't have a notion of "indeterminate truth" – a statement is either true or false in a model. In our approach, we said some identity claims might be indeterminate in an informal sense (like not true in all precisifications). But each *precisification* is itself a classical model. So if we treat the collection of models as epistemic possibilities or as supervaluations, we can say a sentence is super-true if true in all models that respect the vagueness, and super-false if false in all, etc. This is outside first-order logic proper (it's more of a semantic analysis). Formally, our theory might treat "x = y or x ≠ y" as logically valid (we keep classical logic, so tertium non datur holds in each model), but we might refrain from asserting either because we do not know which model (precise assignment) is actual. However, in a given model either x = y or not, no middle allowed [24] [25]. This is consistent with supervaluation semantics that indeed preserves classical validity (Fine 1975 notes that supervaluation yields classical logic at the level of validities, with some caveats [8]). So our formal theory can remain bivalent; the appearance of indeterminacy is a metalinguistic phenomenon of not knowing which model the actual world picks out.

Finally, a word on possible extensions: We have intentionally stayed within first-order logic (no modal operators for necessity or time, no second-order for sets of time-slices). We could extend with a built-in time ordering and even a temporal operator (like a modal "necessarily at future time" operator), but that would leave the core identity framework intact. Similarly, we could add a modal dimension for metaphysical necessity (to talk about trans-world identity vs counterparts) and show that our tags could serve as rigid designators across possible worlds – indeed a tag could be thought of as akin to a trans-world identity key if

one were actualist. But all that is beyond the present scope. Our aim was to ensure that even in the basic first-order setting, we can capture enough structure to answer tricky questions.

In summary, the meta-theory assures us that nothing has gone awry in making these additions: - The system is logically consistent (models exist). - It is a natural extension of FOL= (conservative, presumably no new tautologies except ones involving the new symbols). - Standard metatheorems (completeness, Löwenheim-Skolem, compactness) hold because this is a first-order theory. For instance, if a set of sentences in this theory is finitely satisfiable (each finite subset has a model), then the whole set has a model – compactness. This is important because it implies if some identity scenario cannot be consistently described, it's not because of a mere infinite conjunction or something but a real contradiction in the axioms.

We note a nuance: It's known from philosophic logic that if one tries to incorporate indeterminate identity at the object-language level (like x = y can be "neither true nor false"), you break classical logic quite badly (Evans's argument). We avoid that by not doing it at object level, only at meta-level via multiple models. So indeed, within any fixed model, either x = y or not.

We conclude that our formal groundwork is solid: the theory is robust and doesn't lead to paradoxes or inconsistency, as far as first-order logic can guarantee.

Case Studies

We now illustrate how the proposed framework handles various famous puzzle cases, demonstrating the intuitive outcomes align with reason while avoiding the pitfalls that classical identity alone would face.

Ship of Theseus

Puzzle: Over time, Theseus's ship has all its wooden planks replaced one by one. The original planks are stored. Eventually, a second ship is built from all the original planks. We have: (1) **Replacement**: the ship that results from continuous plank-by-plank replacement of the original, and (2) **Reassembly**: the ship constructed later from the pile of original planks. Which one, if either, is the "real" Ship of Theseus?

Analysis with Provenance and Continuity: In our theory, the original Ship of Theseus at launch (call it S_0 at time t_0) is assigned a unique tag, say $\tau = \alpha$. As planks are gradually replaced, the ship at each later time t_1, t_2, … (S_1, S_2, …) is in continuity with the earlier stages – there's a single continuous world-line, since we imagine the crew is maintaining the same vessel. We enforce an anti-branching policy: as long as there is only one ship in service, there's no branching. By the time all planks have been swapped, the ship in the harbor (Replacement) still carries tag α (numerical identity preserved through continuity and no branching). Now, when the old planks are rebuilt into a ship (Reassembly), that event is a **new origin** – those planks were lying in a warehouse and are now nailed together into a hull again. According to our framework, this Reassembly ship must receive a new tag, say β, reflecting its distinct provenance (the act of reassembly at that later time). It does *not* matter that its material is old; what matters for numerical identity is provenance, and the provenance of Reassembly is entirely different (it was born out of a heap of spare parts at a later date). Therefore, in our theory we can unambiguously say: **Replacement is identical to the original ship (they share tag α), whereas Reassembly is a numerically distinct ship (tag β) despite being made of the original material** [5] . This verdict aligns with practical intuition in many contexts: the ship that

remained in service is considered “the same ship” (legal title, continuous history), whereas the one built from old parts is a *replica* or “the old materials put into a new hull,” not the original ship.

We can further describe the situation using the constitution relation. At the end of the process, Replacement (with new planks) and Reassembly (with old planks) are two different ships. However, one might say Reassembly stands in a certain relation to the original – perhaps it is “the same design and material.” In our terms, we could say at the moment Reassembly is finished, there is an aspect of “having the same parts as the original at launch.” But since the original’s parts got scattered and then regrouped, it’s tricky. What we can say more clearly is: Replacement and Reassembly at the end share no parts; Replacement is constituted by new wood, Reassembly by old wood. There is no direct constitution overlap between Replacement and Reassembly. The interesting constitution relation is between Reassembly and the *original planks*: Reassembly ship ≡_material **PileOfPlanks** (the collection of original planks regarded as a heap). Essentially, Reassembly’s matter is the same as that of the original ship at the start, but via the planks.

However, our framework doesn’t need to dive into that. The key point is it avoids a paradox: classical logic might tempt one to think “if the original is identical to Replacement and also identical to Reassembly, then Replacement = Reassembly, which it isn’t.” We avoid that by denying that the original is identical to Reassembly in the first place – they have different tags, so no identity claim arises. Only Replacement carries the original tag. Thus transitivity of identity isn’t violated [11] because we never asserted the contradictory identities to begin with. The anti-branching axiom prevented us from giving the original tag α to two entities; only one could keep it.

This solution echoes what some authors call the **“best candidate” or “only x and y”** solution (Noonan 1985): one ship at a time. It matches legal practice: you wouldn’t suddenly transfer the ship’s identity to a reconstructed one without any continuity. It also matches engineering intuition: the maintained hull, even if fully replaced, is considered the same vessel (like the proverbial grandfather’s axe that had its handle and head replaced).

One might ask: but isn’t there a sense in which Reassembly *is* the original ship, having all original parts? We can capture any sensible claim there by using the aspect operator. We might say: Reassembly ≡_structure Original (the design and shape are the same, perhaps) or Reassembly ≡_matter Original-at-launch (sharing material kind, though not time). These would be aspectual statements – they do not confer identity, just a kind of equivalence in certain respects. Indeed some philosophers (perhaps motivated by a “four-dimensionalist” view) say both ships are in different senses continuations of the original: one by form, one by matter [26]. Our framework could accommodate that as *different aspects of continuity* (one could define a continuity of form vs continuity of matter). But numerical identity in our strict sense goes with one tag only.

Thus, our calculus cleanly answers the Ship of Theseus: **There is exactly one numerical identity that tracks the ship: the continuously maintained one.** The other is a duplicate with a different origin. This resolves the puzzle without ambiguity and without violating any law of identity [27].

Biological Fission and Fusion

Puzzle: Consider a unicellular organism that divides, or a hypothetical person who undergoes fission (brain split and transplanted into two bodies). Conversely, consider two amoebas that merge into one, or sci-fi scenario of two minds fusing. Who is identical to whom?

Analysis: Our theory insists that identity cannot branch: one origin yields one ongoing individual. So in a fission case, something has to give. We have two modeling choices as discussed: either designate one branch as the "true" continuation (perhaps by some criterion like more brain mass from original, or simply an arbitrary choice), or declare that the original's life ended and two new lives began. Different contexts (legal, biological, personal) might favor one or the other policy. Because our framework makes the policy explicit, we can actually *represent either scenario* in different models without contradiction.

For a concrete example, take the classic philosophical fission: person A's brain is split into two halves, each implanted in a clone body, yielding persons B and C who both have A's memories and personality. According to, say, Derek Parfit, A doesn't survive or at least identity is undefined; according to other proposals (Nozick's closest continuer), maybe whichever is "more continuous" or arbitrary one is A. Let's illustrate both:

- **Closest-continuer policy (arbitrary)**: We decide that B shall carry A's identity. So in the model, we assign $\tau(B) = \tau(A)$ (B gets A's tag) and a new tag to C. Continuity-wise, A's stages up to the split are continuous with B's stages after, but not with C's (maybe break continuity or mark it as a new line). The axiom prevents giving A's tag to C as well. The result: A = B (numerically same person), and $A \neq C$ (C is a new person). This may sound harsh on C, but from a legal or psychological standpoint one might treat C as a "twin" who originates at the fission moment, whereas B is the "original" person continuing. The theory allows this without contradiction. B and C will of course both claim memory of A's life, but only B is A in fact. C's memories are like copied content; we'd say C ≡_psychology A perhaps (same mental template) but not same person (different tag). This scenario would be apt if, for example, we needed to choose who gets A's bank accounts, etc., and we arbitrarily or by convention choose B.

- **Termination-and-duplication policy:** We decide that when A splits, A ceases to exist in the sense of numerical identity (no single one is A), and both B and C are new persons with their own tags (distinct from A's). So in the model, $\tau(A)$ is distinct from $\tau(B)$ and $\tau(C)$, and indeed perhaps A's tag doesn't appear at times after t_split. We set up continuity such that no one is continuous from A's last moment (maybe A's continuity thread terminates), and B and C start fresh (maybe each continuity thread begins at t_split linking to "half of A's state" but since we cannot have a continuity chain from A's single state to two new ones under anti-branching, we simply don't connect A to either – treat it as A died and two similar individuals were born). Then $A \neq B$ and $A \neq C$. The result is: strictly, A did not survive; there are two new people. But we can account for the natural sentiment that "A lives on in some sense in B and C" by saying B and C each share continuity-of-psychology or continuity-of-body with A (but not full identity). We might encode that as aspect relations: e.g., B ≡_personality A and C ≡_personality A (they have the same character or memory up to time of split), or B and C are each constitutionally related to the pre-fission person in some way (a bit tricky because that person no longer exists after fission, but one could treat pre-fission A as constituting part of each new person's history). In any case, neither is the *same tag* as A, so neither is strictly identical. This scenario is akin to Parfit's view that identity is not what matters; A doesn't survive as one, but has two continuers in a looser sense. Our logic handles this by simply not assigning identity and treating any persistence claims beyond fission as indeterminate or false.

Importantly, whichever approach one takes, our framework yields no contradictions. We won't get "B = A and C = A but $B \neq C$" because we never have both B and C sharing A's tag. Either only one shares, or none do. We also avoid violation of Leibniz's Law because we don't assert an identity that would make B and C identical to each other.

From a biological perspective, consider a simpler case: a worm splits into two. Biologists might actually name one fragment as the original (especially if one fragment contains the original head, etc.). If so, that's a policy: the fragment with the head retains the worm's identity, the tailless fragment is a new worm. Or in simpler organisms, maybe they truly treat it as parent ceases, two offspring. Our framework can represent either folk-biological convention by tag assignment.

For **fusion** (two become one), similarly two tags either cease and a new one begins (two amoebas merge into a new amoeba that is neither original strictly – though we might pick one as continuing and the other absorbed), or we pick one as main. Suppose two companies merge. By law, often one is designated the surviving entity (perhaps Company A absorbs Company B, so A's identity continues, B's identity terminates). That corresponds to giving the merged company A's tag, and retiring B's tag. Alternatively, if both cease and a new company C is formed, then neither A's nor B's tag continues – C gets a fresh tag. Both are possible in corporate law; our system can reflect either.

In either case, no formal problems: anti-branching just ensures at most one predecessor tag goes forward.

So our theory explains *why identity cannot branch or merge* in the classical sense: because provenance is unique. When people speak loosely about identity in branching cases, they often end up invoking extra criteria. We make those criteria explicit and keep identity strict. The puzzles are resolved by saying: not both branches can be strictly identical to the original – only one or neither. And not both originals survive as one – either one is subsumed or a new one emerges.

Crucially, our framework also captures the intuition that even if identity fails, something important (continuity, constitution) might hold. For example, after teleportation duplication, even if the original's identity isn't preserved, the duplicate might carry all functional continuity, so in practice you treat it as "same person" for practical purposes except numeric identity. This is captured by continuity-of-psychology vs identity-of-person.

Thus we have a principled way to say: *Identity doesn't branch.* If forced, the semantics marks it indeterminate or picks a branch by convention. This avoids a situation where classical logic would break (classical logic doesn't allow one to assert "A = B and A = C" with B ≠ C, so something had to give; our system simply never asserts the impossible combination).

Destructive Teleportation and Duplication

Scenario: A person enters a teleportation machine. Their body is disintegrated and an exact copy is made elsewhere from new matter (destructive teleportation, like Star Trek's transporter). Is the person who appears at the destination the same person as the one who stepped in? Now consider a malfunction: the original isn't destroyed, so now two copies exist (a duplicate scenario).

Analysis: In the destructive teleportation, the usual intuition is the person survives (since only one exists after, with same memories, etc.), albeit via a very discontinuous process. Our framework can accommodate this by seeing it as a kind of fission where one branch had zero duration (the original's physical continuity was broken). Essentially, original tag could still be given to the new copy at the destination – meaning we treat the teleportation event as not the death of identity but a bizarre transportation. If so, one would have continuity in a loose sense (though spatio-temporal continuity is broken, one could say informational continuity suffices due to the machine's records). This choice corresponds to "transferring the tag" to the

new body. We might justify it by broadening what counts as continuity (maybe a policy that says teleportation output inherits identity if there's no rival copy and it has continuity of mind). Then the copy has the same τ, hence same identity. Our theory yields: original = new (so person survived). Many would like that result.

But note: this required a *stipulation* that the teleporter is identity-preserving. We can also conceive of not granting that – one could argue that the process killed the original and created a clone. In that case, our approach would be to assign a new tag to the reassembled person. Then original ≠ new. The new one is just a different person who just happens to have the same memories (maybe we'd classify them as ≡_psychology but not identical). Philosophers disagree on teleportation; our framework can model either stance by whether continuity is considered preserved or not. If one demands strict physical continuity for identity, then teleportation breaks it, so no continuity → new tag. If one allows continuity of consciousness as sufficient (the pattern continues), then one might allow identity. The key is, whichever we pick, the formalism can handle it.

Now duplication: Suppose the teleporter mistakenly doesn't destroy original, so original remains and a copy emerges. Now we have Person A (original) and Person B (copy), both with the same memories up to teleport. We cannot give them the same tag (that would make them identical and thus one person, which they clearly are not numerically – they can shake hands now). So at most one can have the original tag. This is essentially the fission problem we already did. If we had decided teleportation preserves identity, now two claim it – that's a conflict. In a narrative, usually that's resolved by drama (two claim to be the real one). Our system says: they do not both get to be original. Under anti-branching, one option: treat original A as continuing (keep tag with A, and B gets a new tag). So original A is the "real" person, the other is a duplicate. Alternatively, some might argue the copy is just as good – but logically, one must be designated if any.

Of course, some far-out positions might say in a symmetrical fission, neither is original – identity had "no fact". That is effectively our indeterminate case or termination case as discussed. If it's perfectly symmetric, an impartial stance is to say original identity ceased (tag ended), and both A and B are new. That's also allowed in our framework (especially if we consider fairness: assign new tags to both, original tag is retired). Then indeed "who is original?" has no answer; both are just new, albeit each psychologically continuous.

So in duplication, the framework yields: at most one identical, or none. And the duplicate(s) always get new tags (so they're not identical to original). They may still have continuity of memory, etc., which we capture via continuity/constitution relations. So from the perspective of psychology, we might say "In sense of personality, A ≡_mind B (they are the same mind pattern initially)", but we definitely have A ≠ B in terms of personhood.

This elegantly captures, for example, what goes wrong if one tried naive substitutivity: Suppose original A wrote a will: "I leave my fortune to my future self." If duplication happened, classical identity doesn't tell which of the two gets it (both think they are that future self). In our scheme, if the official policy is original continues as A (no branching, B is new), then clearly A retains rights and B doesn't, even though B remembers writing the will. It might feel unfair to B, but law tends to pick one (like the transporter error scenario could be resolved by deciding one is the legal continuant – which our system formalizes). Or if a judge said "the original died, so neither gets it as continuity-of-person; treat them as separate new persons, the estate is handled as if the person died intestate," that corresponds to the neither branch identity approach.

So the result: *Psychological continuity is handled by our continuity predicate (or an aspect thereof), not by identity.* Numerical identity requires unique origin. Thus, we avoid the classical mistake of saying "they have same memory, so they are the same person," which leads to contradiction if there are two of them. Instead, we say "they share the same memories up to time t (continuity of mind), but they are distinct persons." This solves Parfit's duplication paradox in a straightforward way: memory continuity can branch, identity cannot – so either pick one or declare no identity survived. No contradictions.

Statue and Clay (Material Constitution)

Puzzle: A sculptor forms a statue (call it *Statue*) out of a lump of clay (call it *Clay*). At the moment after creation, Statue and Clay coincide spatially and share all their matter. Yet they seem to differ: Clay existed before (as a lump) whereas Statue did not; Statue has an aesthetic or representational property (being a statue of X) that the lump of clay as mere clay does not; and if we squish the object, Statue ceases to exist (statue is destroyed) but Clay (the clay material) still exists (just deformed). So, are Statue and Clay identical or not? If not, how can two distinct objects be "in the same place at the same time" and made of the same matter? If yes, how do we account for their different properties?

Analysis: Our framework handles this through the **constitution operator** and sorted identity. We represent Statue and Clay as two distinct individuals with distinct tags (different origins). Clay's tag might correspond to when the lump was first formed (say earlier that day, or perhaps it's been around), whereas Statue's tag corresponds to the event of sculpting. Indeed, in many metaphysical accounts, a new object (the statue) comes into being at the moment of sculpting, even though it is made of pre-existing matter. So, $\tau(\text{Statue}) \neq \tau(\text{Clay})$. Thus Statue $\neq$ Clay in the strict numerical identity sense (two different tags). There is no violation of classical identity here: they are simply distinct objects coinciding. This is possible in a many-sorted approach because they can belong to different sortal categories or even if both are individuals, nothing in logic says two distinct individuals can't share all spatial coordinates – that's a contingent scenario.

Now, to explain why we nonetheless talk as if they are "the same thing" in some respects, we invoke the **constitution relation**. We would say: Statue $\equiv$_material Clay (at the creation time and onward, as long as the statue exists). This reads as: the statue and the clay are the same material object or composed of the same matter. More formally, we could say something like: at time t, the *matter of* Statue = the *matter of* Clay (they have identical matter set). That is a way of defining $\equiv$_material. By constitution axioms, this does not imply Statue = Clay, but it does allow us to transfer certain properties. For example, any purely physical property like mass, volume, shape is determined by the matter configuration; since Statue and Clay have the same matter at that moment, they share those properties. So Statue and Clay have the same weight, shape, location, etc. This matches our experience: if you ask "what's that on the table?", you could refer to the statue or the clay – the weight is 2 kg either way, the dimensions are the same.

However, properties that are not purely material can differ. For instance: - **Historical property:** Clay existed at 8 AM, Statue did not (sculpting happened at noon). Our NoMix rule forbids taking "Clay existed at 8 AM" and substituting Clay with Statue, because that property is not purely about material composition at a time, it's about existence over time (which involves identity). So the theory does not conclude "Statue existed at 8 AM" – that remains false (as desired). - **Modal property:** Clay might be able to survive squashing (because clay as a substance can just change shape and still be clay), but Statue cannot (squashing destroys it). Formally, "being squashable" can be phrased as: there is a possible world where the object is squashed but still exists. For Clay, yes (it would just be a lump still). For Statue, no (if it's squashed, that statue as an artwork is gone). Leibniz's Law would say if Statue = Clay, they should share all modal properties, which

leads to a contradiction since one is squashable, the other not [6]. But in our theory, Statue ≠ Clay, so Leibniz's Law doesn't apply between them. We only apply the restricted substitutivity for material properties (being made of clay, having a certain shape), and "is squashable" is arguably a modal/intrinsic property not purely about the matter at one time. So we do not transfer that. Indeed, the constitution view is famous for allowing different modal properties while having same matter [6]. Our aspect framework replicates that: being the same in matter doesn't force same modal attributes, because modal attributes involve identity across possible transformations, which goes beyond pure matter-at-a-time talk.

- **Semantic/aesthetic property:** Statue represents Zeus, Clay does not represent anything (it's just clay). This is a sortal-dependent property: "being a statue of Zeus" makes sense for an art object, not for a lump of clay. We can classify that property under a "statue aspect" or something. Since Clay is not in the aspect of art, we don't attribute that to Clay. Again, no conflict arises because we never said all properties carry over. The NoMix constraint prevents saying "Statue is Zeus-representing, statue = clay in matter, so clay is Zeus-representing." That's invalid because "Zeus-representing" is not a purely material predicate, it's about form and intention. Only statue qualifies.

So, in summary: Statue and Clay are two entities that coincide. **Numerically**, they are distinct (two tags). **Constitutionally**, one (statue) is constituted by the other (clay) at that region of space-time. The constitution relation ≡_material allows us to regard them as "the same object in the material respect," which is why ordinary language often blurs the distinction ("there is one thing on the table, which is both a statue and a piece of clay"). Our theory can say: there is one material thing (the mass of clay) which can be described under two sortals – statue and clay. This is effectively what our aspect logic encodes: inside the material aspect, they collapse to one equivalence class [17], but outside it, they differ.

This solves the classical puzzle by affirming that they are **not** strictly identical (avoiding the direct law of identity conflict), but also explaining why it seems like one object: because of constitution, many predicates will treat them interchangeably. We have blocked the illicit inference from "Statue is (made of) clay" to "Statue = Clay," by understanding the former as an aspectual claim, not an identity claim. Additionally, we enforce "No object can be of two kinds at once unless in an aspect sense" – our sorts or tags ensure the distinction. This addresses the so-called "one too many" problem: it might appear we have two objects where commonsense says there's only one. But actually common sense often implicitly distinguishes the thing as statue vs thing as clay. Our system just makes that explicit. In effect, we endorse a standard metaphysical solution known as "constitution is not identity" (Baker 2000; Johnston 1992) [17] [28], formalized logically.

Hence, the statue/clay puzzle is resolved without mystery: - **No contradiction**: Leibniz's Law is safe because we never asserted identity of the two, and we allow different properties. - **Co-location**: allowed because first-order logic doesn't forbid it; the theory tolerates distinct entities having all their matter overlapped. It's unusual but not inconsistent. We can even express "occupies same region" as a predicate and that holds for them. - **Distinctness**: explained by different provenance: the clay was created in a mixing vat earlier, the statue came into being later. They differ in origin, thus differ in identity. - **Sameness**: explained by constitution: they share the same "stuff," hence everyday observation sees them as one physical object.

This approach is widely regarded as the most coherent for these puzzles, and indeed our formalism was designed to capture it. Our aspect operator is essentially a logical implementation of the idea of an object "considered under a certain description" (in medieval terms, the statue and clay share *materia prima* but not *forma*). We have just ensured that we don't step outside the permitted contexts when doing substitutivity.

Intensional Contexts (Opaque Contexts and Substitutivity)

Puzzle: In intensional contexts like belief reports, necessity, or attribution of names, substituting one identical for another can change the truth of a sentence – seemingly violating Leibniz's Law. Example: Lois believes "Superman can fly" but not "Clark Kent can fly," even though Superman is Clark Kent. Naively, if Superman = Clark, one would substitute and expect Lois believes Clark can fly, which is false – a paradox. Or, "It is necessarily true that 2+2=4, but not necessarily true that the number of planets = 4" even though currently the number of planets is 8 = 2+2+4 maybe a bad example (the classic one: 9 is necessarily >7, the number of planets (which was 9) is not necessarily >7) [29]. How to reconcile these failures of substitutivity with the principle of indiscernibility of identicals?

Analysis: Our framework treats intensional or "typed" contexts as another situation where a straightforward application of identity is blocked by aspect. Specifically, we have a **NoMix constraint** that prevents substitution of co-referential names across certain contexts. We can model a belief context by introducing an aspect like *believed-as* or by treating the objects of belief as tagged differently for the believer.

For instance, Lois Lane has two mental representations: one for "Superman" and one for "Clark Kent," not realizing they refer to the same man. We might formalize that by saying in Lois's belief context, "Superman" and "Clark" are treated as different aspects of the person. We could create an aspect ⟨*Lois's belief*⟩ which distinguishes individuals by how Lois conceptualizes them. It's like saying: Lois's mind assigns two different "tags" to the same real individual, one under the guise "Superman" and one under "Clark Kent." In our system, that can be done by having a separate sort or a separate tagging for mental entities. Another way: treat "Superman" and "Clark" as two different Tags in a subdomain of *guises* that happen to both link to the same real individual externally, but in reasoning about Lois's beliefs, we consider them separate.

The simplest implementation: we say the *belief predicate* or context is intensional and explicitly subject to a constraint: if x and y are identical but have different names, `Believes(Lois, φ(x)) → Believes(Lois, φ(y))` is not an allowed inference unless we know Lois is aware x = y. Essentially, knowledge of identity is required to substitute in belief contexts. This is not an axiom of our base theory, but our architecture prevents making the substitution blindly because `Believes(Lois, P(x))` might be treated as a context where x is taken under a special aspect (like a term in quotes or something).

So we enforce the **NoMix** rule: co-referential names cannot be substituted salva veritate in opaque contexts [30]. This is somewhat an extralogical rule, but in our system it could be encoded by having beliefs be about "tags" or "names" rather than directly about individuals. For example, we could say: Lois believes [the proposition concerning tag T_s (Superman's tag) has property F] and separately [the proposition concerning tag T_c (Clark's tag) …]. If T_s = T_c in reality (they are the same person's tag, if that person is unique), Lois's ignorance means she doesn't realize the tags coincide. One could model that as: in Lois's doxastic state, T_s and T_c are distinct elements (like two different modes of presentation). So within that context, substitutivity fails because the context is effectively sorting things by mode-of-presentation.

To avoid too much complication, we simply state: our theory has room to incorporate these contexts by treating them as aspectual contexts. We can imagine a "belief aspect" K = "according to Lois's belief." Then we could say: **Clark ≡_K Superman** might be false (Lois does not see them as the same), even though Clark = Superman in reality. The NoMix rule definitely prevents going from Clark = Superman to `Lois believes Clark can fly → Lois believes Superman can fly` because the latter

substitution is across contexts. This aligns with the standard resolution in logic: the failure of substitutivity in belief contexts doesn't refute Leibniz's Law properly understood [31], because Leibniz's Law applies to extensional contexts. In intensional contexts, the names or senses matter, which is exactly an "aspect" outside pure identity.

Thus, our formal calculus would treat something like `Believes(Lois, Fly(superman))` vs `Believes(Lois, Fly(clark))` as two different predicates not directly linked by the equality of the referents. We do not have an axiom forcing their equivalence just because superman = clark. We explicitly forbid it unless an aspect says so (like if we had a premise "Lois knows Clark is Superman," that could allow merging the aspects perhaps).

This solves the puzzle by basically adopting the Fregean stance that in propositional attitude contexts, you have reference to senses not just referents [30]. Our system doesn't build a theory of sense, but via the aspect operator, it can simulate having different "identities" for the same object under different descriptions. This ensures substitutivity of identity is only valid in *transparent contexts*. In "opaque" contexts (belief, modality under non-rigid description, etc.), the substitutivity rule is simply not applicable.

This is consistent with standard logic of belief: one might formalize belief not as a first-order relation but via modal operators or some intentional logic. But at a high level, our approach is: the **formal meta-requirements for identity (Leibniz's Law)** remain true in the proper domain (extensional contexts) [31], but we refrain from applying them in contexts they weren't meant for. In the language of our theory, that refraining is done by type-sensitivity or aspect labeling so that "a = b" cannot be freely substituted into a "Believes(Lois, … a …)" context because that context is of a different sort (like involving tags or representations, not the objects directly).

Therefore, intensional opacity is handled not by changing identity's logic, but by **differentiating contexts**. We preserve classical identity rules where appropriate (e.g. outside of Lois's belief, of course Clark and Superman share all properties, including flying ability). But we mark that "Lois's belief about X" is not a property of X in the usual sense – it's a relation involving X under a description. Leibniz's Law does not directly apply [31] (the SEP passage notes that those counterexamples to substitutivity are not counterexamples to Leibniz's Law properly understood, since the content of those sentences is about names or senses).

In sum, our identity calculus provides a systematic way to avoid **substitutivity overgeneration**: only substitute identicals in formulas that do not cross aspect boundaries. Belief contexts are one such boundary. The logical rigor we introduced with sorts and aspects thus naturally prevents one from making the mistake that leads to paradox. Lois believing "Superman flies" is formalized in a way that does not allow the replacement of "Superman" with "Clark" even though it's the same person, because it's not the same role in her belief system.

This addresses not just belief but also things like reference in fiction, or cross-world identity in modal logic if you treat it similarly (counterpart vs identity could be seen as a sort difference as well). However, modal necessity is often handled by rigid vs non-rigid designators. Our tag approach ties to that: a rigid designator (like a tag or a proper name fixed to an origin) would preserve identity across worlds, whereas a description like "the number of planets" is non-rigid, hence identity substitution fails in necessary contexts. That example (9 vs number of planets) in the SEP lines [30] is exactly about a non-rigid designator – something

our theory would treat perhaps as a function of time or world. We won't elaborate further, but the resources are there.

To conclude case studies: each puzzle – Ship of Theseus, fission/fusion, teleportation, statue/clay, and intensional substitutivity – is solved by one of our three identity relations: - **Numerical identity via tags** ensures uniqueness and prevents multiple occupancy of one identity. - **Continuity** explains persistence and lets us choose non-arbitrary outcomes for branching (or explicitly admit indeterminacy). - **Constitution/ aspect** allows co-location and intensional distinctions without confusion, by introducing controlled equivalences for specific respects.

These solutions are achieved *while preserving* the classical ideals (we never made identity intransitive or relative in the formal sense; identity is absolute but we supplement it, and we never abandoned Leibniz's Law for the identity predicate itself – we only avoided misapplying it).

Comparison with Competing Views

It will be helpful to contrast our approach with other major theories of identity that attempt to handle these puzzles:

Classical Identity (orthodox view): The standard view in logic and much of philosophy is just to use one absolute identity relation with no further structure. Puzzles are then addressed by either refining language (talk of "same F" vs "same G") or by philosophical argument that the puzzles aren't real problems of identity per se (some, like Lewis, claim all the fuss is about something else, not identity [32]). For example, classical theorists might say in Ship of Theseus: "identity has no fact in such cases or it's a matter of arbitrary choice – but identity itself is trivial" (Lewis 1986; see quote in SEP) [32]. Our view differs in that we do not dismiss these puzzles as ill-posed or extraneous; instead, we *add* structure to identity to solve them. We retain the classical formal laws (reflexivity, etc.), so we agree with classical identity on what those laws are, but we *relocate their role*: the formal laws remain as meta-constraints on the base identity relation (tags), while the "criteria" for identity are given by provenance, etc., not just by an unanalyzable primitive. Thus, where the classical view would say "x = y or not, and if hard cases, oh well, maybe it's indeterminate but logic can't handle that," we say "x = y has a determinate truth if tags align, and if not, false – and any indeterminacy is in whether the tags align, which is a factual vagueness." So we see ourselves as extending, not overthrowing, the classical view. We ensure that all the formal constraints of identity that classical logic cherishes hold true of our numerical identity relation (tag equality). For everyday situations where identity is clear, our approach collapses to the classical one: tags are like trivial names for individuals, continuity and constitution relations either trivial or not invoked, so nothing changes. But in tricky cases, we have extra expressive power to avoid paradox. In that sense, we could call our theory a *conservative refinement* of the classical picture: we don't change any classical theorems, we just solve things that were underdetermined by adding postulates. So we get the best of classical logic (no weird logical consequences like relative identity does) while addressing its limitations.

Relative Identity (Geach 1967 and others): Relative identity theory denies that there is a single absolute identity relation that holds across all sortals. Instead, it says identity statements are only meaningful when qualified by a category: "a is the same F as b" might be true while "a is the same G as b" is false [7]. For example, proponents of relative identity might say the Clay and Statue are the same material object but not the same artifact – and they'd treat that as literally "identity" relative to a concept. In relative identity logic, one typically does not assume the standard transitivity or substitutivity across categories. While this sounds

somewhat like our approach, we actually *avoid* the main controversial commitment of relative identity. We still have an absolute identity (tags) in our ontology. We aren't saying identity is fundamentally sort-relative; we're saying there are other equivalence relations (like continuity, constitution) that are *not* identity but can mimic it in limited domains [17]. Geach's view, more radically, might say absolute identity is meaningless and only same-F is meaningful. We reject that: we have absolute identity (two things having the same provenance tag *are one thing*, period). We thus preserve what many see as essential: the idea that counting and quantification require an absolute identity at base (Quine 1964 argued that without absolute identity, we can't even count objects properly) [33]. Our theory agrees: tags allow us to count individuals (if two share a tag, they count as one, otherwise two) [34]. So we consider relative identity approaches to be unnecessary. Instead of scrapping absolute identity, we augment it with relational predicates for "same F in aspect." So in technical terms, we keep identity as a logical constant (with standard axioms), rather than treating it as a mere family of equivalence relations. We do, however, achieve some of the aims of relative identity: we can say "the statue and clay are the same material thing" effectively with ≡_material, which is what a relative identity theorist would say "same lump" but not "same statue." The difference is we never confuse this with actual identity, and we maintain an overarching view that one object (tag) can manifest as different sortal instances. Our formal NoMix constraint ensures that, unlike naive relative identity, we don't accidentally violate core properties. Relative identity theory has been criticized for leading to logical issues (some argue it's incoherent because it denies that "the same F and the same G" implies same thing, etc.). Our approach is immune: we don't actually allow contradictions because we isolate contexts. We might be described as having something like "relative indiscernibility" rather than relative identity.

Counterpart Theory (Lewis 1968, 1986 for modal; sometimes applied to temporal identity as well): David Lewis suggested that instead of literal trans-world identity, an object in one world is a different object in another world that is a "counterpart" – extremely similar. While that's about modality, a similar idea can apply to time (the stage theory: an object at one time is not literally identical to itself at another time but is a counterpart (a later stage) in the same world). The counterpart approach solves certain puzzles by not having one enduring entity, thus avoiding transitivity issues by saying strictly, an object's stage now and its stage later are distinct entities connected by continuity. Our approach actually doesn't require going to that extreme. We allow one to speak of an object being truly identical over time (because it has one tag) – we are more in line with an endurantist view that an object is wholly present at different times, but we articulate what that means via continuity. We didn't need to break identity into disjoint stages – though our semantics can be seen as having stages (the (x,t) pairs) but those are connected strongly by the continuity relation. In effect, we capture some benefits of stage theory (like handling fission by noting multiple counterparts without one identity) but we do so within a single-world first-order theory and still being able to say, if no branching, that it's "the same object." Counterpart theory in modal logic helps avoid certain rigid designator issues; in temporal logic, counterpart-like thinking helps avoid "temporary identity" paradoxes by saying x-at-t is a different entity from x-at-t′. We didn't go that far; we treat x at different times as the same if the tag is same. But we implemented a mild version: we have an equivalence relation connecting those times. So one can reinterpret our continuity relation as a kind of counterpart relation along the time axis, but one that has been tamed by the anti-branching rule to act like true identity when unique. In any case, counterpart theory often introduces a lot of machinery and sometimes complicates counting (if you say strictly x-at-t1 and x-at-t2 are different, how many objects? We still say one tag thus one object). We avoid reliance on counterpart theory for ordinary intra-world identity. We confine counterpart-like ideas to how we might interpret modally different origins (we haven't deeply done modals here, but presumably each tag is like a *thisness* – cross-world, either that tag exists or not, giving a fixed trans-world identity or not depending on worldview, but we can remain neutral).

Essentialist or Aristotelian Approaches: Some theories say that identity conditions are given by essential properties or sortal kinds. For example, Wiggins (2001) argues every object falls under a sortal (like person, tree, etc.) that determines its persistence conditions (what changes it can survive) and thus its identity through time. Kripke (1980) introduced ideas like the necessity of origin – which is an essentialist criterion, basically an object cannot have a different origin and still be the same, and certain properties like biological species membership are essential to identity. Our framework is quite compatible with these notions: in fact, provenance tags embody the "necessity of origin" principle inherently (if you try to imagine the same tag object with a different origin, that's a contradiction – each tag is a particular origin event). So we have built in something Kripkean: if x had been born of different parents, it wouldn't be x because the tag would differ (that matches Kripke's intuition exactly). The difference is we formalize it logically rather than just philosophically. We however stop short of requiring additional essences for identity beyond origin – we put continuity as a separate relation, not baked into identity (some essentialists would try to incorporate maybe soul or form as an identity principle – we separate that out as constitution aspect). We treat typical essentialist truths (like "this table could not have been made of ice instead of this wood" – an origin essentialism statement) as something that can be captured either by tags (if it's about origin) or by constraints on constitution (like maybe the table's *wooden composition* is essential, meaning you cannot replace all matter with steel and call it the same table – that we would capture by saying the continuity would break if such a drastic change occurred, or that being wood is part of its "essence aspect" and if that's lost, the aspect identity stops). So essentially, we relocate essentialist conditions not to identity directly but to continuity and constitution rules. That is, rather than say "x's identity is defined by having property P," we say "losing P means continuity is broken or a new tag is needed or an aspect identity no longer holds." This is more flexible and avoids the often debated problem of what exactly is essential. We can accommodate variety: e.g. some might say personhood is essential to an individual – in our system, that could be reflected by saying if an individual ceased to be a person (went into persistent vegetative state?), maybe continuity as the *same person* would end (it becomes just a human organism, different aspect). That's an application of our aspect operator rather than identity per se. So we can echo what sortal essentialists say, but we do it with logic of aspects, not by complicating the identity relation itself.

So relative to essentialist approaches: we share the spirit that identity is not just a blank slate – things have nature. But we choose to express natures in the constitution sort, leaving identity itself simple and uniform. This avoids fights about whether identity can be contingent or not – in our system, identity is necessary if true (because a tag either refers to same in all possible worlds or not at all – though we didn't formalize modal logic, presumably if you treat tags rigidly, x = y implies necessarily x = y, a property we maintain [35]). Distinctness is also necessary if we assume origin essentialism globally (which we effectively do: if $\tau(x) \neq \tau(y)$ then in no possible world would they be one, assuming tags are tied to actual origins; we might get into haecceitism vs anti-haecceitism debate with that, see below).

Haecceitism (thisness) vs Anti-Haecceitism: A concern might be that our tags are effectively *haecceities*, like individual essences or thisnesses (Adams 1979). We respond: the tags in our theory are not mysterious primitive "identity entities," they are more like logical proxies for a concrete causal origin or event. We aren't reifying a transcendent property "being identical to X" – we are assigning a parameter that corresponds to something like "the event of X's coming into existence." In implementation, a tag could be (Time, Location, microstructure) of origin etc. We treat it as atomic for generality, but it's conceptually grounded. Therefore, while we do have a form of haecceitism (in that each individual carries a primitive identifier in the model), we align with a moderate view: we do **not** allow two distinct individuals that share all qualitative properties to nevertheless differ (our tags are anchored in origins, and if qualitatively everything is same but they come from different origin events, then yes they are distinct – that's a Black's spheres scenario (Black 1952)

where two indiscernible spheres differ only by being two – our theory would assign different tags because presumably two separate creation events or positions). That's actually a haecceitist position: we do allow that scenario, because nothing in our logic forces two qualitatively identical things to be identical. But that's just classical extensionality anyway. We just give a name (tag) to that difference. We can explain the difference by different provenance. So to someone who worries we've smuggled in primitive thisness: we can say it's not primitive in an epistemic sense – it corresponds to something observable in principle (like separate causal origin or distinct history). We are fine with a mild haecceitism: the world could have had two qualitatively identical objects swapped, which our tags would distinguish (that's akin to a permutation of tags making a different possible scenario – yes, we allow that, which is basically embracing haecceitism as the denial of purely qualitative supervenience of identity). But since tags are logical tools, this doesn't inflate our ontology with ghostly properties, it just acknowledges that identity facts might not reduce to qualitative facts – a stance some (like Adams, Kaplan) also take. If one is anti-haecceitist, one could add a constraint that tags must correspond to qualitative unique properties or something, but we don't do that.

Perdurantism (4D theory) vs Endurantism (3D): Our continuity approach is neutral enough to accommodate either viewpoint. If one is perdurantist, one could interpret each tag as identifying a 4D worm, and continuity simply enumerates its temporal parts. Our anti-branching then aligns with the notion that a worm can't branch into two worms and still be one worm (the Y-shape would be two worms sharing a segment). Stage theorists (a version of perdurantism) would treat each time-slice as a separate entity and "identity" as an I-predicate linking slices. Actually, our continuity relation essentially functions like the "genidentity" relation connecting slices into a worm (Carnap used that term). So a perdurantist could see our individuals as 4D sums and the continuity relation as trivial (each individual trivially continues as itself) or see our individuals as slices and the continuity as the identity of worms. Endurantists could see each tag as a continuing substance that wholly persists – which is our primary talk. The difference is mainly philosophical interpretation: our logic doesn't commit either way – it just provides an equivalence on stage pairs. So we in effect unify some aspects of both: we can talk like an endurantist (x is literally the same as x at t2 if same tag) but handle branching like a perdurantist by saying in branching you have two distinct continuants because you can't have one tag follow both – similar to how perdurantists would say fission yields two overlapping worms that shared the first part. They might say the original person was actually two people all along sharing early stages. We don't quite say that, but we say at fission moment nature forces a distinction.

The above comparisons show: our theory is something of a hybrid approach that keeps the best parts of classical absolute identity and essentialist rigor (clear principles for identity), while introducing something akin to the fine-grained distinctions of relative identity or counterpart theory but without their more problematic implications.

Finally, one particular **application: theology**. Some proposals, notably in Christian theology, regarding the Trinity (Father, Son, Holy Spirit being one God but not one person) have used ideas akin to relative identity or constitution. Our framework can naturally express that scenario as an application (we do this in the next section). We mention this because one might suspect our theory was motivated by such a theological concern (the mention of "relational-first Catholic theology"). However, our stance is the theory was developed independently as a general identity calculus, and it just so happens it can clarify such theological claims without special adjustment. That stands in contrast to relative identity theory, which Geach partly motivated by the Trinity, or other ad hoc proposals made solely to solve the Trinity puzzle (some theological models impose mysterious identity-like relations). Our theory covers that case as just another instance of constitution relations: the persons of the Trinity are distinct (different provenance, one might even

poetically say the Father unbegotten, Son begotten, etc., which are different origin relations – those could correspond to distinct tags), yet they are of one divine essence (which could be an aspect in which they are all one). That is exactly a constitution-like situation (distinct persons, one substance). We can implement that with our logic by saying Father ≠ Son ≠ Spirit (different tags, thus distinct persons), but Father ≡_divine Son ≡_divine Spirit (all the same substance in aspect of divinity). And the NoMix rule prevents concluding "Father = Son" or collapse, just like statue/clay. This lines up with the classical Nicene formulation (three hypostases, one ousia – essentially three personal subjects constituted by one divine nature). We thus succeed in providing a logically coherent model of the Trinity that doesn't require exotic relative identities (unlike e.g. van Inwagen's relative identity defense). We mention this because it showcases the versatility of our calculus: it can handle philosophical theology without special-case adjustments.

Applications beyond Metaphysics

While our theory was formulated with deep metaphysical puzzles in mind, it has very practical analogues in other fields. The ideas of provenance, continuity, and constitution correspond to common notions in information systems, law, and science:

- **Information Systems (Databases, Version Control):** In databases, especially distributed ones, one often assigns objects a unique identifier (like a GUID) that remains the same through updates – this is exactly the idea of a provenance tag. For example, a customer record might have a primary key; even if the customer's details change, that key identifies the same conceptual entity. If duplicates are found (two records for the same real person), deduplication procedures effectively discover a single provenance for what was thought to be two and merge them (assigning one identity). Conversely, if one record was mistakenly representing two people (like two individuals share a name), you split it and give distinct identities. Our provenance identity can model **data deduplication** by saying "two entries with different keys are distinct unless we unify keys." It also models **versioning**: a new version of a file inherits the original's tag if it's an edit, or gets a new tag if it's a separate file. Many version control systems (like Git) treat content hashes as identity markers – if a file is copied and modified, is it same or new? The user or system uses heuristics (continuity of file path, etc.) to decide. We could simulate such policies.

Also, in knowledge graphs or object-oriented contexts, there's the problem of object identity vs equality by properties. Our approach clarifies that – it's the classic thing where two objects have all the same attribute values but are still distinct instances (because different IDs). That resonates with the idea of haecceity in our tags. So in computing, our theory appears as unique IDs plus reference integrity.

- **Legal Chain-of-Title and Traceability:** In law, the concept of a piece of property maintaining identity through transfers or modifications is crucial. For instance, the Ship of Theseus problem actually appears in maritime law (if you replace enough parts, is it the same registered ship or a new one?). Laws often stipulate thresholds or conventions (the "hull identification number" or what counts as the same vehicle after repairs – often the chassis number serves as identity tag). Our framework formalizes such practice: the anti-branching policy in law might say if two halves of a company split, maybe the original's rights go with the branch retaining the original name or charter (this is done in corporate law: one company is designated the successor). This exactly matches picking which gets the original tag.

In intellectual property, chain-of-title means tracking provenance: who created a work (tag), and then how rights flow by continuity (assignments, etc.). If a work is derivative but sufficiently new, it might get a new identity (like a new tag for a derivative work vs a copy is same tag context). Similarly, in supply chain traceability, each product has a lot code or serial number (tag) to trace back to origin (for recalls, etc.), and assembly continuity records which part went into which, such that identity of a composite can be resolved (the Ship of Theseus in manufacturing: if 60% of a car is replaced by spare parts, is it same car for legal VIN? Usually yes if chassis unchanged). Our formal continuity and provenance can encode those rules so that queries like "is this the same item that was certified earlier?" can be answered by checking tags and continuity chain.

- **Biology and Manufacturing Traceability:** Biologists track cell lineages (provenance of cells through division – which is basically an identity tree). Our anti-branching resonates with the idea that you can't have one cell's identity go to two daughters equally – one might call them new. (In biology, usually both daughters are new – original cell doesn't survive fission). This is a policy choice as well. Genetic lineage tracing uses tags (like certain atoms or markers) to see continuity. In manufacturing, serial numbers and batch numbers serve to maintain continuity of identity through transformations (like raw material batch to final product). If two batches merge, either one loses identity or you assign a new batch number for mixture. That parallels our fusion story.

- **Databases and Temporal Persistence:** The continuity predicate essentially is like an object's history. Temporal databases assign stable IDs to objects and track their state over time, which is exactly separating identity from state. The anti-branching policy might correspond to whether the database allows splitting one entity's history into two – usually not without creating a new entry. Our logic could be implemented in a temporal database to enforce consistency (no object can have two different values at same time unless it's considered two objects).

- **Software Engineering – Assemblies and Versions:** Consider software artifacts composed of modules. If you replace a module, is it the same system? In semantic versioning, a patch version means it's the "same" API (continuity preserved), whereas a major version increment might signal not backward compatible – arguably a new identity. We can model these decisions: e.g. an API might have a provenance identity (the project) and continuity until a big change where they might deliberately break continuity (rename or give new ID). Also, layered software (like a virtual machine on hardware – the VM is constituted by the hardware's resources, but not identical – similar to statue/clay scenario). Our aspect operator clarifies that: the VM and the host machine share the same bits in memory (material aspect) but the VM is a distinct entity (it can be moved to another host – then constitution changes, host doesn't follow). This is analogous to soul and body in philosophy or any dual-aspect system (we could say a mind is constituted by a brain – not identical but intimately related. That's a bit like dualism moderated by constitution relation).

- **Engineering objects**: In engineering, BOM (bill of materials) and configuration management revolve around identity of assemblies. Our theory can formalize "part X at time t is the same part as at time t0 if same serial, provided it hasn't been replaced." If replaced, new serial implies new identity. "Ship of Theseus" in aerospace: an airplane's hull number is fixed; if you rebuild it from new parts gradually, it's still the same plane by reg number – that's exactly using a tag. But if you reconstruct one from scraps, it would get a new registration. So our rules are already informally applied. We just give them a logical foundation.

Turning to the **theology application** (as promised, briefly): Our framework wasn't designed for it, but it "transparently supports" it. In Christian theology, there's the claim "God is one being (one substance) in three persons." If we interpret "being/substance" as the aspect of divinity (what they are), and "person" as individual identity (who they are), we can model Father, Son, Holy Spirit as three distinct individuals (tags) that exist in the same eternal context, and define a constitution relation ≡_DivineNature that holds among them – effectively saying they share the same divine nature. We also add that each is not constituted with any other besides in that single aspect (so no transitivity issues beyond them; it's an equivalence between them that they all are the same God in aspect). This yields formally: Father ≠ Son ≠ Spirit (distinct persons), but Father ≡_God Son, Son ≡_God Spirit, and by symmetry and transitivity of that equivalence in aspect, Father ≡_God Spirit (they are all one God). Any property that is divine (omnipotence, eternity) is ascribed via the divine aspect, so it's shared (NoMix ensures if Father has property of being almighty, and that is a divine-nature property, then Son and Spirit also have it) – which matches doctrine (they are co-equal in divine attributes). Meanwhile, person-specific properties (like being the Father of the Son, or being born as Jesus in time) are not transferred across that aspect – those are personal or relational properties, not essential to divinity, hence NoMix forbids, say, "Father generates Son" from implying "Father generates Himself" or "Son generates someone," we keep them separate. This aligns with the idea that persons are distinct in relations. Essentially, we get a consistent model of Trinity that avoids both polytheism (because one aspect shows unity) and modalism (we do not collapse the persons – distinct tags means distinct personal identity).

We mention this as an "illustration rather than constraint": It shows our calculus is flexible and general, not that we were motivated by theology in designing core axioms. The core axioms were motivated by secular puzzles; the theology just provides a neat confirmation that our structure (one substance, three persons is exactly one tag-aspect linking three tags in our terms) was general enough to handle something even as tricky as Trinity, which relative identity theorists have struggled with.

Thus, the architecture proves useful widely: - It gives databases and software a rigorous handle on identity vs equality. - It models legal and historical identity questions clearly. - It matches scientific practice of tracking lineage and composition.

One could foresee implementing these ideas in, say, a **proof assistant or data management tool**. For example, a knowledge base might use our sorted logic to avoid conflating an entity's identity with its properties. Or a digital provenance system might formalize anti-branching to ensure uniqueness of identity through merges/splits.

Objections and Replies

We anticipate several objections to our approach, which we address in turn:

Objection 1: "Provenance tags are just a dressed-up form of haecceity or primitive thisness. Haven't you just assumed what identity is (a tag) rather than explained it? It seems like you collapse into saying each thing has a primitive identity property (the tag), which is no better than just saying each thing is identical to itself by a primitive."

Reply: The provenance tag is indeed a kind of primitive label in the formalism, but its significance comes from what it represents: an object's origin or causal lineage. We are not merely saying "because x has a

primitive tag α and y has α, therefore x = y" as a tautology; we are asserting that **origin** is the criterion for numerical oneness. This is a substantive claim: it means if two things originate from the same event or cause, they are the same thing (and conversely distinct origins entail distinct things). This links to philosophical principles like Kripke's necessity of origin (which we embrace). The tag is a logical device to implement that principle without getting entangled in natural language. It's not an extra metaphysical entity or mysterious haecceity beyond the object; it's more akin to an event ID or a placeholder for "the process that produced this." In practice, one could often replace tags with descriptions of origins (e.g. "the clay that was shaped by sculptor at 2pm on Monday" could serve as a unique description for the statue's tag). We keep them atomic for formal simplicity, but they are conceptually grounded. So, while yes, we have something like a haecceity, it's not brute or unanalyzable – it's analyzable as an origin reference. This is different from just saying "each thing has an irreducible thisness." Instead, each thing has an origin, and that origin is crucial and often empirically or historically accessible. The benefit of the tag approach is that we can rigorously track identity through changes and duplicates by referencing origin, rather than juggling a vague notion of "the same object." It's not circular because we reduce identity to an externally given relation (same origin), rather than leaving it as a self-referential property. Also, in many fields (e.g. archival science, art, data management), provenance is an accepted concrete concept – by leveraging it, we demystify identity. So, far from being a cheat, this is an explanatory gain: you ask "what makes this the same object over time?" we answer "it came from the same source, see tag." If someone insists that's just a brute fact, we counter that any theory will have to have some primitive notion of identity or unity; at least ours ties it to a natural concept (origin) instead of leaving it completely featureless.

To draw an analogy: saying each electron has a "thisness" doesn't help, but saying each particle has a worldline or trajectory from a creation event is informative and often measurable. We incline to the latter – tags parallel worldlines labelled by creation.

Objection 2: "Continuity with an explicit policy sounds ad hoc or conventional. Isn't identity supposed to be an objective fact? If your identity over time can depend on a 'policy' (like how we choose in fission cases), doesn't that make identity arbitrary or subjective? One might say you are smuggling in a non-identity criterion (like survival or best continuer) and calling it identity. Is that legitimate?"

Reply: We maintain that in the vast majority of cases, continuity – the persistence of an object – is an objective matter, given the nature of the object. The "policy" aspect only comes into play in exceptional scenarios where the usual criteria yield ambiguity (like symmetric fission). In everyday life, objects don't branch; people don't split into two, ships don't duplicate spontaneously. So continuity is straightforward to track (the same person has one continuous life, etc.). In those cases, our anti-branching policy doesn't need to intervene; identity follows the single path of continuity.

Now, in thought experiments or rare cases (like identical twins from one zygote, or Parfit's brain split, etc.), we admit that there is an element of convention or decision in saying which resulting individual is "the same" as the original, or whether to declare the original ended. But this is a virtue of our framework, not a vice: it acknowledges that some identity questions **have no determinate empirical fact** to resolve them and that forcing a definite answer would be artificial. We prefer to explicitly incorporate that decision as a parameter (policy) rather than let it invisibly bias results. Consider the alternative: those who say identity is perfectly factual either have to pick an answer arbitrarily or deny the scenario is possible. By making policy explicit, we ensure the theory does not inadvertently assert something unwarranted.

The policy selection is not utterly whimsical either; it is typically guided by the context's purpose. For example, law or practical concerns may designate one branch as the successor because that is useful or fair (as in companies merging or splitting). In personal identity debates, someone might pick the branch with more psychological continuity as "really you" – that's a policy guided by a reasonable principle. Our framework can accommodate whichever principle one endorses, by building it into the continuity relation. If one feels no branch has a claim, one can represent that as no identity continuation (both are new).

So yes, identity in extreme cases can be conventional – but that reflects how people actually handle those cases (we see this in legal fictions or in Sci-fi where sometimes they shrug and say "both are the original in their own way" – which in our model means original tag retired, two new tags). The crucial point is, once a policy is fixed in a given interpretation (a given model), the identity relation regains objectivity: it becomes an equivalence relation that either holds or not, no fuzziness. The policy doesn't "smuggle subjectivity into identity" in the final analysis; it picks one admissible candidate relation. After that, everything is back to being factual. And where no policy is fixed because maybe none is appropriate, our theory gracefully allows indeterminacy rather than forcing a false precision. That is arguably the correct stance: better to say "it's indeterminate whether A = B" (because scenario is ill-defined under classical identity) than to assert something and get contradictions or violate intuition.

In short, identity is usually factual; when it's not, our theory says so and requires a decision or treats it as indeterminate. This is actually similar to how identity is treated in some contexts – for instance, legal identity of an entity after radical change may require a court decision (policy). We've just formalized that notion. This transparency about policy is an advantage – it prevents hidden assumptions. And it doesn't undermine the concept of identity because in normal conditions the policy is implicit (e.g. we assume no duplication, so policy doesn't come into play).

One can also view our continuity relation as generalizing the "closest continuer" theory academically – but whereas that theory was an informal proposal for personal identity (and rightly criticized for being arbitrary or teleological), we incorporate it in a rigorous way: we specify it as a parameter, not a built-in metaphysical law. Thus, we can explore consequences of different choices and it's clear which is chosen. That is a methodological improvement, not a weakening of identity.

Objection 3: "Your aspect/constitution operator sounds like just saying *relative identity* in another form. You say statue and clay are 'the same in aspect X but not strictly identical.' How is that fundamentally different from someone like Geach saying 'the statue and clay are the same lump of matter but not the same artifact'? It seems you've just renamed relative identity as constitution or aspectual sameness. So have you actually avoided the pitfalls of relative identity or just relabeled them?"

Reply: The critical difference lies in how strictly we control the use of the aspectual equivalence and how we preserve an underlying absolute identity framework. In Geach's relative identity approach, the idea is that "x is the same F as y" can hold while "x is the same G as y" does not, and there is no overarching identity relation at all. That leads to potential logical problems (like one can't count how many things there are without referencing a sort, etc.). In our framework, however, we **do** have an overarching identity relation (tag equality). Statue ≠ Clay absolutely. That means we never blur the line that they are two distinct entities in general. The aspect relation is an additional structure that lives, so to speak, on top of that absolute distinction. We ensure that the aspect relation does not undermine the absolute distinction by enforcing NoMix and non-compositionality.

Thus, while it might sound similar (“same in this respect, different in that”), our formalism doesn’t license problematic inferences that relative identity might. For example, relative identity logicians sometimes face issues like: if a and b are the same F and b and c are the same G, can a and c be compared? In our system, all such comparisons reduce to underlying tags if you go outside aspect contexts, so we don’t end up with logical ambiguity at the top level. The aspect operator cannot be freely iterated to cause contradictions because it’s not transitive across different aspects or contexts. Each ≡_K is a separate equivalence relation domain. We disallow any inference from one domain to another (NoMix). So the collapse that critics fear in relative identity (like it being trivial or inconsistent) is avoided. Our approach might be seen as a sort of *typed identity*: each aspect context has an identity-like relation, but the types ensure they never conflict with each other or with global identity.

Another way to see the contrast: relative identity theories often deny that statements like “x = y” (unqualified) have meaning or are ever true except under a sort. We do not deny that – we very much allow “x = y” and consider it the primary identity claim (with all laws intact). We simply also allow “x ≡_K y” for various K, which relative identity would interpret as “x = y as K.” But because we still have a notion of x and y being simply identical or not, we avoid the scenario where identity becomes entirely relative. We have the cake (absolute identity) and eat it too (also a notion of sameness in some respects).

So indeed, constitution is not identity [17] – we adopt that slogan literally as an axiom: being made of the same matter doesn’t make you identical, just like relative identity folks say. But we don’t generalize that to say identity is always relative. We keep constitution as a special, explicitly limited case, used only in specific contexts (like coincident objects or intensional contexts). It's a controlled tool, not a free-for-all reinterpretation of '='.

Finally, perhaps the question implies “aren’t you just doing what relative identity people wanted to do, except with extra formal caution?” Our answer is maybe yes, in a sense we achieve what relative identity advocates desired (the ability to talk about same-this but not same-that) without sacrificing the classical logic of identity. We feel that’s a strength – it vindicates the useful insight of relative identity (that sameness needs sortal context in many cases) while avoiding its excess (denying absolute identity or making logic non-standard).

One could say we have a two-tier identity: a strict numerical identity and a secondary equivalence notion for aspects. This layered approach is fundamentally different from a single-tier but many-sorted identity of Geach’s approach. Therefore, our aspect operator isn’t merely a renaming of relative identity; it’s a different animal that coexists with absolute identity. This is why, for instance, puzzles like “many-one identity” (one thing constituting many or vice versa) don’t cause contradictions for us – we simply say many things can constitute one thing in different aspects, and that’s fine, whereas a pure relative identity theory might struggle to express that clearly.

Objection 4 (if any, though three were requested): Perhaps some might question the **complexity** of the system – isn't it too baroque to be plausible? But since not explicitly asked, I'll stick to the three.

In summary, none of these objections undermines the viability of the theory. Tags may seem like an extra assumption, but they encapsulate origin which is a powerful explanatory base [8] . Continuity with policy may seem to introduce convention, but only where identity is already problematically indeterminate; otherwise it captures what we already implicitly do in practice. And the constitution relation is carefully constrained to avoid the paradoxes of relative identity, giving us the ability to talk about “same in one sense,

different in another" rigorously [17]. These features collectively solve puzzles while preserving classical logic's virtues, which is a strong balance to strike.

Conclusion

We have presented a first-order theory of identity that *goes beyond* the classical minimalist conception by providing an explanatory framework for what makes something one and the same. The theory divides the labor among three notions: - **Provenance-based numerical identity:** The strict, classical identity relation is preserved, but we explicate it via provenance tags – two entities are one when they share an origin. This satisfies all formal laws of identity while giving them content (an object's identity is *grounded* in its coming-to-be). We thus meet the formal requirements (reflexive, etc.) within each domain just as classical logic demands, but we no longer treat identity as metaphysically featureless; we pin it to a cause or source. - **Continuity for persistence:** Instead of forcing identity through time to do all the work, we use a continuity relation (like a world-line) to handle an object's persistence conditions. This relation adheres to an anti-branching principle, reflecting the insight that identity doesn't split. By doing so, we cleanly handle issues of fission and fusion without violating identity's transitivity – any would-be problematic case is either resolved by stipulating a single continuation or flagged as indeterminate rather than breaking the law. The gains are evident in all diachronic puzzles: we can talk of something surviving changes or ending, with continuity rather than absolute identity doing the heavy lifting when changes are complex. - **Constitution for aspectual equivalence:** To address cases of coincident entities and intensional contexts, we introduced a constitution/aspect operator that captures sameness in a given respect (material, form, context) without equating the entities outright. This allows us to say, in effect, X and Y are "two aspects of one reality" or "the same in essence, different in person" in various scenarios – statue vs clay, or the persons of the Trinity, or an object under different descriptions – without sliding into contradiction. The formal safeguards (non-transitivity across individuals, limited substitution) ensure this behaves well.

By assigning these different roles to different relations, we avoid overburdening the pure identity predicate with tasks it cannot perform (like handling change or intensional contexts) and avoid misusing it in places it leads to paradox. The classical meta-requirements for identity – reflexivity, symmetry, transitivity, Leibniz's Law in its proper form – are all **preserved where they belong**: namely, for the relation of sharing a provenance (and within each aspect context appropriately). We have not thrown out the rulebook of classical logic; we've extended it in a conservative way to incorporate needed distinctions. The result is that we can solve the standard "stress test" puzzles of identity: duplications, fissions, fusions, Ships of Theseus, constitution problems, substitutivity puzzles, etc., in a uniform formal system. No ad hoc exceptions or paraconsistent maneuvers are required – everything is handled with classical reasoning plus our extra predicates.

This division of explanatory labor – provenance for numerical identity, continuity for persistence, constitution for qualified sameness – yields clarity. Problems that were paradoxical under a simple "just same or not" view dissolve because we recognize they were asking different questions (about origin, survival, or constitution) and our theory can answer each on its own terms. For example, the Ship of Theseus: "Is it the same ship?" becomes two questions in our analysis – does it have the same tag (origin)? Yes, the maintained one does. Does the rebuilt one? No, new tag. Hence only one is numerically identical. No contradiction emerges. Or statue and clay: "Are they one or two?" – two numerically (different tags), one materially (constitution). No contradiction – just precision. This pattern repeats across puzzles.

We believe this framework aligns well not only with metaphysical intuition but also with practices in other fields as shown. It provides a **calculus of identity** that can be implemented in formal ontology or computer science (where tracking identity vs change is crucial). One could even imagine mechanizing these rules in a **proof assistant** or database system to automatically detect identity conditions or inconsistencies. Because it is a first-order theory, in principle it could be encoded in a theorem prover. All our axioms are first-order; we could attempt to verify scenarios (for instance, verify that certain transformations do or don't preserve identity given the axioms). The conservativity ensures this enrichment doesn't break existing reasoning about equality – a comforting thought for integration with existing logical systems.

Looking ahead, there are modest extensions that could enrich the framework without upsetting the core results. One could add a **modal dimension** – talk about possible worlds – and by treating tags as rigid designators, likely maintain that if x = y then necessarily x = y (which we already expect) [35], and handle counterpart-like reasoning for cases where continuity or constitution differ across worlds (this might help formalize essence vs accident precisely). We could also integrate a **temporal logic** or mereology (to talk of parts continuity) – but again, those would layer on top of what we've built, not requiring changes to identity itself.

The core thesis, to reiterate, is that identity is not mysterious or paradoxical once we articulate *what makes something one*. By going beyond "= as primitive" and giving it structure through provenance, continuity, and constitution, we preserve logic and solve paradoxes. In metaphysics, this clarifies debates on personal identity, artifact identity, and so on. In law and tech, it matches how identities are tracked. And intriguingly, in philosophy of religion (an application we mentioned), it naturally models unity and multiplicity without special pleading.

In conclusion, the classical notion that identity is "utterly simple" – just a = a – is true in a formal sense, but it left many practical and metaphysical questions unanswered. Our theory shows that we can keep that formal simplicity (in each well-defined domain) while providing a rich explanatory apparatus around it. The result is a disciplined, many-sorted logic of identity that is equipped to handle complexity without lapsing into inconsistency. This is the "identity calculus" that philosophers, computer scientists, lawyers, or theologians could use when ordinary '=' is too coarse.

By reassigning the explanatory burden to provenance (for why there is one thing and not two), to continuity (for how one thing persists), and to constitution (for in what sense two things can be one), we have not only solved classical puzzles but also provided insight into everyday reasoning about identity. The work ahead could involve formalizing this in a proof assistant to check consistency and perhaps including **modest modal or temporal enrichments** to capture necessity of origin or time-indexed identity fully – but we expect those to integrate smoothly, given the careful design. The core framework stands firm and intact after all these examinations, promising a powerful tool for anyone grappling with the concept of identity in a rigorous way.

Bibliography

Adams, R. M. 1979. "Primitive Thisness and Primitive Identity." *Journal of Philosophy* 76: 5–26.

Aristotle. 1984. *Complete Works of Aristotle*, ed. J. Barnes. Princeton: Princeton University Press.

Avicenna (Ibn Sina). 2005. *The Metaphysics of The Healing* (al-Shifa'). Translated by M. Marmura. Provo: Brigham Young University Press.

Aquinas, Thomas. 1981. *Summa Theologiae*. Translated by the Fathers of the English Dominican Province. Westminster, MD: Christian Classics.

Baker, Lynne Rudder. 1997. "Why Constitution Is Not Identity." *Journal of Philosophy* 94 (12): 599–621.

Brower, Jeffrey E., and Michael C. Rea. 2005. "Material Constitution and the Trinity." *Faith and Philosophy* 22 (1): 57–76.

Fine, Kit. 1975. "Vagueness, Truth and Logic." *Synthese* 30 (3-4): 265–300.

Geach, Peter T. 1967. "Identity." *Review of Metaphysics* 21: 3–12.

Hawley, Katherine. 2001. *How Things Persist*. Oxford: Clarendon Press.

Koslicki, Kathrin. 2008. *The Structure of Objects*. Oxford: Oxford University Press.

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

Lewis, David. 1986. *On the Plurality of Worlds*. Oxford: Blackwell.

Merricks, Trenton. 2001. *Objects and Persons*. Oxford: Clarendon Press.

Noonan, Harold W., and Ben Curtis. 2018. *Identity (2nd edition)*. London: Polity Press.

Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Clarendon Press.

Sider, Theodore. 2001. *Four-Dimensionalism: An Ontology of Persistence and Time*. Oxford: Clarendon Press.

Thomasson, Amie L. 2007. *Ordinary Objects*. New York: Oxford University Press.

van Inwagen, Peter. 1990. *Material Beings*. Ithaca: Cornell University Press.

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

Williamson, Timothy. 1994. *Vagueness*. London: Routledge.

1 3 9 10 15 18 21 22 29 30 31 33 Identity (Stanford Encyclopedia of Philosophy)
https://plato.stanford.edu/entries/identity/

2 4 5 6 7 11 12 13 14 16 17 19 23 26 27 28 32 Identity Over Time (Stanford Encyclopedia of Philosophy)
https://plato.stanford.edu/entries/identity-time/

8 [PDF] and It Can't Be Supervalued Either Jerry Fodor and Ernie Lepore ...
https://ruccs.rutgers.edu/images/tech_rpt/tr-19/evaluate19.pdf

20 24 25 Vagueness (Stanford Encyclopedia of Philosophy)
https://plato.stanford.edu/entries/vagueness/

34 Persistence in Time | Internet Encyclopedia of Philosophy
https://iep.utm.edu/per-time/

35 Relative Identity (Stanford Encyclopedia of Philosophy/Summer 2017 Edition)
https://plato.stanford.edu/archives/sum2017/entries/identity-relative/