ChatGPT

Licensed Identity from Causes: A Participation-Friendly Replacement for Classical Identity with a PEI Surface

Abstract

Identity is reconceived not as a primitive logical constant but as a licensed verdict that becomes true only when deeper causal and ontological conditions are met. In this framework, an identity claim is "earned" by satisfying truth-makers in provenance (common origin and continuous history) and integrity (persistence conditions), rather than assumed outright. Meanwhile, classical equality ($=$) survives intact as a conservative shadow within purely extensional discourse, reproducing standard logic where it applies. This paper makes three contributions. **First**, it eliminates primitive identity in favor of a *licensed identity* relation (`Id`) defined from causal provenance and constitutive profile: sameness of origin, sameness of constitutive nature-set, and continuity constraints are required for `Id(s,t)` to hold. **Second**, it introduces a readable surface language (a PEI-style presentation) with **typed predication** and **constitution** symbols that clarify aspect and category, preventing category errors. **Third**, it provides a single "participation-friendly" calculus that uniformly handles puzzles of material constitution (Ship of Theseus, Lumpl vs. Goliath), fission and fusion cases, personal identity through change, and even theological doctrines (Trinity, Incarnation) without ad hoc exceptions. The system conservatively extends classical first-order logic with equality – no theorem about "=" in extensional contexts is lost or gained – yet it reframes what identity *is* and how it functions by deriving it from more fundamental ontological relationships [1] [2]. In sum, identity becomes the *endpoint* of explanation, not the start, allowing the logic to resolve longstanding paradoxes while remaining compatible with orthodox metaphysical and theological principles.

1. Introduction and Claim

1.1 The Problem: Classical vs. problematic identity contexts

Classic first-order logic treats identity as an absolute, primitive relation: reflexive, symmetric, transitive, and governed by Leibniz's Law (the indiscernibility of identicals). In ordinary extensional contexts this works brilliantly – one can replace a with b in any truth-functional statement if $a=b$, preserving truth. Indeed, identity so understood is "trivial" in the sense that each thing is identical only to itself and to nothing else [3]. However, this orthodox view of identity strains under a variety of well-known puzzles. It tells us *that* each thing is one and the same with itself, but not *why* or *how* we determine when something remains **the same** under changing conditions [4] [5]. For many philosophically interesting situations, the simple $a = b$ framework underdetermines the facts:

- **Material Constitution:** Cases like **the Ship of Theseus** and the statue vs. clay (**Lumpl and Goliath**) expose a gap in classical identity. The formal laws alone do not determine which future individual is "really" identical to the original when two equally continuous candidates emerge (e.g. the ship rebuilt from original planks vs. the ship gradually repaired) [6] [7]. In the statue/clay scenario, classical logic would demand that if a lump of clay L and a statue S share all the same atoms at a time, then perhaps $L=S$; yet L and S have different properties (the statue has an aesthetic form, the clay does not; the clay could survive being squashed, the statue could not). Classical identity cannot hold without violating Leibniz's Law, leading to paradox if we naïvely identify them.

- **Persistence and Fission:** Thought experiments in personal identity (for example, a person's brain splitting into two bodies, or an amoeba dividing) raise the famous fission problem: one original becomes two continuers. By classical logic, two distinct future persons cannot both be identical to the one past person (by transitivity and symmetry), yet each future individual may seem in some sense continuous with the original. Classical identity offers no degrees or structured relations here – one must arbitrarily pick one continuer as "really" the same or insist that identity miraculously survives as a one-to-two relation, which standard logic disallows. Similar issues occur in fusion (two things merging): classical identity cannot represent "two became one" except by denying that two ever existed or that one is genuinely new. In short, when confronted with **branching or merging histories**, classical identity lacks the resources to describe what is going on, except to declare identity lost or violated.

- **Intensional and Modal Contexts:** In intensional contexts (belief, knowledge, modal necessity, etc.), substituting identicals can lead to falsehood. For instance, Saul **Kripke** observed that Lois Lane can believe "Superman can fly" while not believing "Clark Kent can fly," even though Clark Kent = Superman [8] [9]. If $x=y$, classical logic allows replacing x with y in any context, but here that yields a truth ("Lois believes Superman flies") turning into a falsehood ("Lois believes Clark flies"). Leibniz's Law breaks down in opaque contexts. Similar problems arise with necessary identity: $a = b$ might be true but not necessarily true (if a and b are names for the same person but it's a contingent fact they are that person). The classical theory has to chalk these up to context sensitivity or ban substitution in those contexts, essentially an ad hoc restriction rather than an explanation.

- **Mixed-sortal Predication and Theology:** In metaphysics and theology, we encounter statements that defy a simplistic reading of "is" as identity. For example, in Christian doctrine *"Jesus is God"* and *"Jesus is a man"* are both true, but it would be heretical and illogical to conclude "God = a man" by transitivity. Similarly, "the Father is God" and "the Son is God" are true, yet the Father is not the Son. Clearly, the copula "is" in these cases is not expressing a plain identity between individual supposits, but something like predication or a qualified identity (often explained as "identity in nature" or *suppositum* to nature relationship). Classical first-order logic with a single category of object and a monolithic identity relation is ill-suited to articulate these distinctions – one either gets contradictions (if "is" is identity) or must refuse to formalize such statements. As theologians put it, the **doctrinal grammar** of Trinity and Incarnation requires that some "is" statements be parsed in a special way (e.g. "is God" meaning "possesses the divine nature" rather than absolute identity) [10] [11]. The classical identity apparatus has no native way to represent one subject having two natures, or one nature belonging to multiple persons, without breaking the identity relation or adding kludgy side conditions.

In summary, **classical identity shines in purely extensional contexts** like mathematics and everyday tangible comparisons, where it provides a clear criterion of when two names or descriptions refer to the same entity. But it **struggles in contexts of constitution, change, intensionality, and multi-aspect entities**. The formal laws themselves don't solve these puzzles – they only highlight the contradictions when we try to force a yes-or-no identity verdict. As one introduction to these issues notes, there is a difference between the formal constraints on identity and "the truth-makers or grounds that make identity claims true in real situations" [12] [5]. Classical logic leaves those truth-makers unspecified. We are left asking: on what *basis* do we say "these two occurrences are one and the same thing"? And could rethinking that basis allow a more fine-grained, yet principled, logic of identity that *avoids* the above paradoxes?

1.2 The Proposal: Licensed identity from causes (LFI-π) with a PEI surface

This paper proposes a new paradigm in which identity is *derived* rather than assumed: a relation that must be **licensed** by satisfying specific criteria of sameness in origin and composition. We call this approach **licensed identity from causes**, abbreviated **LFI-π** (the π evoking *provenance* or origin tags). The idea builds on the Provenance–Esse–Integrity (PEI) framework introduced in earlier work [13], but here streamlined and made more "participation-friendly." In practical terms, we introduce a formal identity predicate `Id(x,y)` *defined* by a conjunction of conditions capturing what it means for x and y to be the "same being" in a robust sense. Specifically, **x and y are identified (`Id(x,y)`) only when there exists some origin tag π such that**:

- **Provenance:** x and y share the same originating cause or history π. We write $R_π(x,y)$ for "x and y have the same provenance tag π." Intuitively, π could be thought of as a unique "birth identifier" or world-line: for example, the tag of the particular lump of clay in the statue scenario, or the tag representing a specific continuous life in a fission scenario. If two entities do not stem from the same causal origin, they cannot be collapsed into one; provenance is a necessary condition for identity [14] [15].

- **Constitutive Profile:** x and y have the exact same set of constitutive natures or forms. We write $\mathrm{Con}(x)$ for the set of all substantial natures under which x exists (its constitutive profile), and require $\mathrm{Con}(x) = \mathrm{Con}(y)$ for identity. For most mundane cases, an entity has only one substantial nature (e.g. Fido's nature set is {Canine}), so this condition is trivial. But in cases of **multi-natured entities** (such as the Incarnate Christ with both divine and human nature), this condition is crucial: it ensures we only assert identity if the entities do not differ in what they fundamentally *are*. If one object is a lump of clay (nature = {clay}) and another is a statue (nature = {clay, statue-form}), their constitutive nature sets differ; they cannot be identified. This prevents conflating things like a person with their nature or a statue with its material when those differ in definition. Identity thus demands sameness not just of matter but of form-profile – a stronger condition than classical equality which cares nothing about *kind*.

- **Integrity (Persistence):** x and y must satisfy the integrity conditions appropriate to their kind, ensuring continuous existence without disruptive breaks. We denote this $I(x,y)$: roughly, y is the "continuous continuation" of x (and vice versa) according to the relevant criteria for that sort of entity. Integrity covers things like spatio-temporal continuity, functional or psychological continuity, etc., depending on what kind of thing we are describing. It plugs in the *sortal-specific persistence conditions* that philosophers like Wiggins argued are needed for identity judgments [16]. For example, a living organism might require biological continuity (no death and replacement) to count as the same; a ship might require a certain continuity of form and material. In our framework, different **integrity policies** can be chosen for different sorts (more on this in Section 5.2), but the key is: if x "dies" or is destroyed and y is a recreation or replica, then $I(x,y)$ fails, and identity is not licensed. Integrity thus rules out "resurrections" or gap-filled identities unless one explicitly allows those via the policy. It also means apparent same-origin cases that violate continuity (like branching) will fail identity (addressed next).

- **No Branching (Anti-Multiplicity):** Even if x and y share an origin π, that origin must not have split into two simultaneous continuations. We require a **no-branching condition** $NB_π(x,y)$: along the history tagged by π, there is only one continuing entity at a time. If an origin line "forks," identity cannot hold between the products of the fork (at most one can retain the original tag as its identity). This captures the intuition that an individual cannot split into two numerically identical individuals – if splitting occurs, at least one of the resulting entities is a new individual with a new tag [17]. In practice, if at time t_0 we have one entity A with tag π, and at a later time t_1 two entities B and C both claim tag π (a branch), the rule will force that at most one truly keeps π. The other must get a different provenance (say a new tag π') marking it as distinct. This no-branching clause cleanly handles fission puzzles: it "downgrades" what would have been a classical identity into a weaker sibling relationship when branching occurs [18] [17].

All together, we **define the identity of x and y via a chosen tag π**:

$$ \textsf{Id}(x,y) \;:\iff\; \exists \pi \Big(R_{\pi}(x,y)\ \wedge\ \mathrm{Con}(x)=\mathrm{Con}(y)\ \wedge\ I(x,y)\ \wedge\ NB_{\pi}(x,y)\Big). $$

In words: *x is identical to y if and only if there is some common origin π such that x and y share π, have the same constitutive natures, satisfy the integrity/persistence conditions, and π's line does not branch between x and y.* Only when all these conditions are **jointly** met do we "license" the assertion that x and y are one and the same entity. Identity thus becomes a *derived verdict* – an endpoint reached when certain explanatory conditions hold, rather than an unconditioned primitive. As a slogan: **identity comes *from* causes** (and continuities), not from sheer logical fiat.

Critically, this scheme does *not* throw away classical equality. Instead, it **contains classical equality as a special case**: in purely extensional contexts where those extra conditions are either irrelevant or uniformly satisfied, `Id(x,y)` behaves exactly like $x=y$ does in standard logic. We effectively **preserve Leibniz's Law and the classical equivalence relation on a safe subdomain** of discourse, but we refine it in contexts where intensional distinctions or multi-category distinctions matter. For example, in a simple mathematical domain or a set of indivisible atoms, one can stipulate that each entity has a unique tag and trivial integrity (they never branch or change), so our `Id` reduces to "having the same tag" which is just a one-to-one marker – essentially replicating $=$ [19] [20]. The new identity relation is an **equivalence relation** (reflexive, symmetric, transitive) on each sort of entity when considered under extensional, non-problematic conditions [21] [22]. In those conditions, it also allows unrestricted substitution just like classical identity (we will formalize "licensed substitution" rules in Section 3.3). Thus, **LFI-π is a conservative extension of first-order logic with =**: any formula purely about the "=" relation in a context where identity is unproblematic will still hold if and only if it was a classical logical truth [23] [24].

What we gain, however, is that **outside** those neat conditions, the system does *not* allow certain inferences or identifications that lead to paradox. Identity becomes something you have to *prove* by showing all identity conditions are met, rather than an automatic assumption. If any condition fails – say, two things have the same origin but diverged into different forms, or they share form and continuity but come from different sources – then you simply cannot assert $x=y$ in our system (or if you do, it will be false). The calculus will block inferences that would be illicit. For instance, we won't be able to substitute "the statue" for “the clay” in a law-like statement that doesn’t tolerate their difference in properties, because the logic knows they are not truly identical but only constitutionally related (more on that in a moment). Similarly, we won’t conclude that the two Persons of a fission scenario are identical to the original, because NB fails and identity is not licensed past the split – we instead get something like *“shared origin up to time t”* but no further [17] [25]. In intensional contexts, we will mark certain predicates or contexts as *extensional positions* where substitution is allowed versus *intensional positions* where one must explicitly show provenance unity to substitute. This disciplined approach internalizes what used to be ad hoc side-conditions (like “don’t substitute in opaque contexts”) into the logic’s basic licensing rules.

To implement this approach in a user-friendly way, we introduce a **PEI-style surface language** (call it **PEI-L**) that sits atop the formal calculus. “PEI” stands for *Provenance, Essence (esse), Integrity*, reflecting the three pillars of the approach [1]. The surface language uses intuitive symbols for different kinds of statements: - We write $s \in N$ to assert **predication**, meaning “supposit s is of nature N” (for example, “Socrates ∈ Human” reads as “Socrates is human” – a true predication of a nature to a person). - We write $s \diamond N$ (a diamond or similar sign) to assert **constitution**, meaning s is constituted under nature N (i.e. N is in the constitutive set Con(s)). For many cases $s \in N$ and $s \diamond N$ will coincide (a person has one nature). But if a being has multiple natures (like Christ who $∈$ Divine and $∈$ Human, and thus $◇$ Divine and $◇$ Human both hold for him), the constitution symbol marks that explicitly. Constitution can be seen as “sameness in aspect” without full identity [26] [27]. - We write $R_π$ as a binary relation symbol for **same provenance/origin**. So $R_π(s,t)$ is usually abbreviated by something like $s ≡_π t$ in the literature [28], indicating s and t have identical origin tag π. Provenance is provided with *indices* or tags; π itself might carry information (e.g. temporal span, origin type). - The ordinary equality symbol “=” will appear in the surface language only in places where it is safe – essentially as a **macro** for `Id(s,t)` when used in a purely extensional context. We will ensure in the formal rules that any occurrence of “=” in user-level discourse either expands to a licensed identity claim or is confined to a context (like pure mathematics or set-membership statements) where it behaves classically. In other words, “=” is not banned, but when a reader sees $a = b$ in this paper, they can trust that either (a) it’s in a fragment of discourse where identity is trivial (and thus means the same as classical equality), or (b) it’s a shorthand, with the understanding that behind the scenes $Id(a,b)$ was proven under the required conditions.

By combining these notations, we can neatly distinguish statements that classical logic would otherwise conflate. For example, **“the statue is (the same material object as) the clay”** can be represented as `Statue ◇ ClayStuff` – a constitution claim – rather than `Statue = Clay`. The latter is simply false in our system (they fail Con(.) equality and provenance differs), whereas the former is true and meaningful: it says the statue and the lump *share the same matter or substance stuff* without asserting they are numerically identical [27]. This nuance allows us to avoid Leibniz’s Law problems; e.g., from `Statue ◇ Clay` we do **not** infer that every property of the statue is a property of the clay, because `◇` is not identity and has its own guarded rules (it’s a kind of “aspectual sameness” relation, **non-transitive** and bound by sortal categories). Only when something is truly `Id` with something else can you substitute freely, and our logic makes sure that happens only when it should.

To summarize the proposal: **LFI-π (Licensed-from-cause Identity)** reframes identity as a *judgment dependent on prior conditions*. It “licenses” claims of equality only when same cause, same form, and continuous integrity concur without branching. In all other cases, we refrain from asserting identity, using weaker relations like predication or constitution to express partial unity. This approach preserves the successful parts of classical identity (its use in extensional reasoning and mathematics) but extends the logical language to be *typed* and *guarded* in ways that handle the problematic cases uniformly. Identity is no longer the fundamental explainer of why a thing is one; rather, identity emerges as a *verdict* after explaining the origin and persistence of the thing [2] [29]. As a result, the logic naturally avoids the classic paradoxes without needing case-by-case exceptions. Section 3 will formally develop the core calculus of LFI-π, and Section 4 will illustrate the surface syntax (PEI-L) that makes it accessible to work with.

1.3 Payoff: One calculus across philosophy and theology (without special pleading)

What does this new framework buy us? The payoff is a **single calculus** that can accommodate a wide range of identity puzzles and cross-category inferences in a principled way. Instead of having separate rules or entirely different notions of identity for metaphysical problems versus theological ones, we get one coherent system that addresses them all by the same foundational strategy. A few highlights:

- **Material Constitution Puzzles:** In our calculus, the **Ship of Theseus** problem is resolved cleanly. We represent the original ship with a unique tag π. As planks are gradually replaced, the ship at each moment keeps tag π (since continuity and integrity hold). If someone later reconstructs a ship from the old planks, that reconstructed ship has a **different** provenance tag (say π_2) – it doesn't matter that it has the same planks as the original at an earlier time; it was reassembled later, ergo a different origin. Therefore, the rebuilt-from-old-parts ship is **not** $\equiv_\pi$ to the continuous ship (which stayed in service), and our system will not license saying they are identical. They will simply be two distinct ships that share a lot of similar material. Classical logic struggled here because it has no place to register "coming from the same stock of planks" as distinct from "is the same ship." We do: one ship gets tag π, the other gets a new tag, so Id fails between them [30]. At the same time, our logic can articulate a **sibling relationship**: we might say the two ships are "origin-equivalent up to a point" or that they stand in a *continuity* relation before the branching. But after the branch, one ship is π, the other π_2 – numerically distinct. This aligns with common-sense: we want to say the refurbished ship (with new planks gradually installed) *is* Theseus's ship, whereas the reassembled one from cast-off planks is a replica. Our system captures that intuition rigorously: identity requires no branching and single provenance, which the replica lacks with respect to the original [7] [17].

- In the **Lumpl vs. Goliath** (clay vs. statue) case, classical logic forces an unhappy choice: either say they are identical (and then somehow explain away the property differences), or say they are two coinciding objects (raising questions about how two objects can share all matter). Our approach says: The lump of clay and the statue have **different constitutive profiles** (the lump's nature is just a lump of clay, the statue's nature is a lump of clay organized in a statue form – effectively an extra formal property). Hence $\mathrm{Con}(Lump)\neq\mathrm{Con}(Statue)$, so by definition `Id` cannot hold. They are not identical – which matches Leibniz's Law since, for instance, the statue has the property "being a statue" that the lump lacks. However, we don't need to treat them as utterly unrelated either. We can say the statue is constituted by the clay: $Statue \diamond ClayMatter$` is a true statement capturing their intimate connection [27]. This constitution relation lets us talk about *"the same material object in one sense but not absolutely the same entity"*. It **blocks improper substitution**: from "Statue is fragile" we cannot infer "Clay is fragile" because our logic knows "Statue" and "Clay" are not fully identical, only materially the same in aspect. So the puzzle is resolved by acknowledging a layered unity (material continuity with formal distinction) instead of a flawed identity claim. This is very much in line with the notion some philosophers have that co-located objects share matter but not identity [31]. Our single identity calculus handles it by *licensing identity only at the right level* (the clay today is the same clay as the clay yesterday – that we allow, same tag; but the statue today is not the same entity as the lump before sculpting – different form set, so no identity).

- **Fission and Fusion:** When one entity splits into two (like a person fissioning, or an amoeba dividing), classical identity can't say "A = B and A = C" without inconsistency, so one might be forced to deny real identity or say identity is lost. Our system handles this by the tag and branching rule: only one of the outputs can keep the original tag π. The other must get a new tag (call it π'), indicating it's effectively a "new" entity from the split moment forward [17]. Before the split, of course, there was just A with tag π. After, we have B (tag π) and C (tag π'). So $Id(A,B)$ was true up until t (before splitting) and remains true after between the appropriate stages, whereas $Id(A,C)$ fails at the moment of split (because NB fails—two offshoots). The calculus can still express a relation between B and C, perhaps a **sibling-of** relation or simply that they were one before t (we formalize that as sharing π up to a point, which can be expressed if needed). But strictly, beyond that point, identity is not carried forward to both – one is the "continuation" (A's identity persists in B), the other is a new offshoot. This matches common judgments: if you split a worm, perhaps you'd say one half is the original and the other is a new worm, depending on context (or you might say two new worms – either way, you don't maintain a single identity in two tracks). **Fusion** (two become one) is handled dually: if two tags α and β fuse into one entity D, then D must pick one tag as primary (say α) or get a wholly new tag γ signifying a new entity created by fusion [32] [33]. Either way, at most one original identity carries on; the other is subsumed or terminated. Again, this avoids contradictions. Classical logic said nothing about how to treat "two become one"; in ours, provenance bookkeeping yields a consistent story. Notably, in both fission and fusion, our logic yields a **"weaker" relation in place of identity** at the troublesome point – something like *was a part of the same continuing history* or *contributed to the constitution of*. This reflects reality better: in fission cases, we often say **the two resulting persons were one person** (past tense), or in fusion, **the single resulting entity was two entities** (past tense), without asserting they are identical going forward. Our calculus can represent those truths without inconsistency. In short, it delivers rigorous rules for these scenarios where classical logic had only silence or brute force (indeed, one author notes that without further stipulation, "Leibniz's law doesn't directly tell you what to do when one becomes two or two becomes one" [33] – we have filled that gap with explicit rules).

- **Personal Identity through Change:** The debate in personal identity often centers on which continuity (bodily, psychological, memory-based, etc.) is the "true" criterion for identity over time. Our framework doesn't legislate one ultimate answer; instead, it provides a logical container to **plug in different integrity policies** and see their consequences. If one adopts a **biological continuity** integrity condition, then `I(x,y)` will demand unbroken functioning of the organism between x and y for identity. If one adopts a **psychological continuity** criterion (say, for persons, memory and personality preservation), then $I(x,y)$ will encode that. The rest of the calculus stays the same – you still need same origin and no branching. Thus, our system can express rival theories as different *interpretations of I*. It also allows nuanced scenarios: for example, one might formalize cases of gradual replacement (like teletransportation in sci-fi) by adjusting what counts as integrity (maybe spatio-temporal continuity is broken but if the information continuity is preserved and we're generous, we count that as integrity – or maybe not, depending on the policy). Rather than building one answer into the logic, we give a framework to **articulate the differences**. The main point is that, whichever policy one chooses, it must be stated explicitly, and identity claims will be objectively *true or false* depending on whether that policy's conditions were met [34] [35]. This approach encourages clarity: one cannot just say "maybe they're kind of the same person in some respects" without formalizing which respects. You either meet the declared I condition or you don't. If you don't, identity is not licensed – though you might still describe other relations (e.g. “x and y share memories” as a weaker fact). This positions the framework as a tool for exploring identity in various domains – from law (chain of custody identity), to computer science (version control and object identity), to psychology (self over time) – by encoding domain-specific integrity constraints without changing the logical rules.

- **Theological Doctrines:** Perhaps most strikingly, LFI-π applies to theological puzzles **without special ad hoc rules**. The logic naturally maps onto the classical Christian understanding that, for example, in the Trinity **there is one divine nature but three distinct persons**. How so? In our system, one would say: Father, Son, and Holy Spirit are three distinct supposits (entities of sort S). Each of them has $\mathrm{Con}(Person) = {\textit{Divine Nature}}$, meaning each is constituted by the one same divine nature. So we assert $Father ∈ Divine, Father ◇ Divine$, and similarly for Son and Spirit [36] [37]. However, the Persons do **not** share any common origin tag: by doctrine, the Father is ungenerated (no origin), the Son is generated from the Father, the Spirit proceeds from Father and Son [37]. We denote this, for instance, by assigning tags: $\pi(Father) = U$ (think “uncaused”), $ \pi(Son)=F$ (“from Father”), $\pi(Spirit)=FS$ (“from Father and Son”) [38]. These tags are all distinct. Therefore, for any two distinct Persons, say Father and Son, there is no single π such that $R_π(Father, Son)$ holds – their provenance differ (one has none, one has source in Father). Thus `Id(Father, Son)` **does not hold**, and similarly for any pair of Persons. Our calculus thereby validates “Father is not the Son, etc.” as a fundamental fact, not merely a primitive exception, but as a consequence of their different origin relations [39]. At the same time, we *can* formally express “Father, Son, Spirit are consubstantial (share one nature)” by the predicate $P(s,Divine)$ or by saying $\mathrm{Con}(Father) = \mathrm{Con}(Son) = {\textit{Divine}}$ (and a special rule that multiple distinct individuals can all constitute the same single nature kind – essentially that the nature itself isn't an individual that would violate identity). The logic even allows us to articulate the doctrine of **inseparable operations**: any action A that is of the divine nature (say, creating the universe, or performing a miracle) can be represented as an extensional predicate on the nature (or as something like $A \in \textit{DivineOperation}$). Given each Person $P ◇ Divine`, the calculus will support an inference that if a Divine-operation is true of the Divine nature, it is true “of” each Person in that nature (each Person will be associated with that act) [40]. Thus we can derive statements like “The Father creates, the Son creates, the Spirit creates” from “Creation is an act of God” without implying three separate acts – it will effectively attribute one and the same act to all three (because in the extensional predicate view, the operation belongs to the one nature). This models the ancient formula *“the external works of the Trinity are undivided”* [40]. All this is achieved **without** needing to change or break our logical rules for the theological case – we simply interpret the tags and types appropriately. In classical logic, representing “three Persons, one God” is awkward at best (it usually requires multi-sorted approaches or non-standard identity like “relative identity”). Our system manages it with standard identity (`Id`) but enforced only where appropriate: since no two Persons meet the identity criteria, we never equate them, yet we can still use logical reasoning to relate them via the shared nature.

- In the doctrine of **Incarnation**, the claim is that the eternal Word (Second Person of the Trinity) took on a human nature as Jesus of Nazareth, such that Jesus Christ is one person (hypostasis) in two natures, divine and human. How to formalize *"one person, two natures"* without confusion? In our framework, it is natural: let W be the divine Person (the Son/Word) and N the human nature that began at the Incarnation (through Mary). We have initially $\mathrm{Con}(W) = {\textit{Divine}}$. At the Incarnation event, W assumes a human nature, so thereafter $\mathrm{Con}(W) = {\textit{Divine}, \textit{Human}}$. No new supposit is created; it is one and the same W (the tag for the Word continues unbranched through this assuming). Thus $Id(\textit{Pre-Incarnate Word}, \textit{Incarnate Word})$ holds because it's literally the same person continuing (same tag, now with expanded Con set, which requires that integrity conditions allow an accretion of nature – we can stipulate that's allowed for divine persons). Meanwhile, Jesus as a man is just W considered under the human nature aspect. We can use $P(Jesus, Human)$ to say *"Jesus is human"* and $P(Jesus, Divine)$ to say *"Jesus is divine"*. The calculus will prevent us from equating the human nature with the divine nature (different sorts) or Jesus's humanity with his divinity – those are in his Con set, not objects to identify. So one *supposit* has two entries in Con. From this we can analyze statements: *"Jesus is God"* in our logic becomes "Jesus (the Person) has the Divine nature," i.e. $Jesus ∈ Divine$. *"Jesus is man"* becomes $Jesus ∈ Human$. Both are true. Yet if someone mistakenly tried to say "God = Man" on the basis of that, it would fail because "God" and "Man" are natures, not supposits, and our system never allows a supposit-to-nature identity (`Id` requires both arguments be same sort, either two supposits or two natures, not a mix). Thus the dreaded confusion of natures is prevented by sorted typing alone [26]. Moreover, **communicatio idiomatum** (the "communication of properties" between Christ's two natures) can be handled elegantly with aspect controls. For example, we can formalize the statement *"God died on the cross"* in a careful way: The subject "God" here refers to the Person (the Word), and "died on the cross" is a predicate true of him *according to the human nature* (since God in His divine nature cannot die, but the Person as man can die). In our calculus, one would prove *"W (who is Divine) has the property of mortality at time t"* by noting $W ◇ Human` and that mortality is a property of the human nature aspect; thanks to our rules, this doesn't imply anything absurd like "the Divine nature itself died" – it is contextualized. As the Council of Chalcedon taught, one must attribute things *"theologically correctly": the same one person has both divine and human predicates, keeping each nature's properties distinct* [10]. Our formal system mirrors this: it allows *personal predication* across nature boundaries but within aspect, and stops any illegitimate substitution that would swap an entity's nature-bound property into the other nature's context. In Section 7.2 we will walk through how formulas can express classically paradoxical statements ("Mary is the Mother of God", "God suffered and died") in a way that is logically consistent – indeed, **provable** – under our approach [41]. Importantly, none of this required adding "exceptions" to logic for theology; we simply use the same rules of LFI-π with an appropriate model where, e.g., no single π spans multiple divine Persons, and where having two natures is allowed for one tag (without branching).

- **Sacramental Metaphysics (Transubstantiation):** As a final example of breadth, consider the Catholic doctrine of the Eucharist, where at consecration the substance of bread is converted into the substance of Christ's Body, while accidents (appearances of bread) remain [42] [43]. Philosophically, one might ask: is the consecrated host *identical* to the bread that was on the altar moments before? The appearances suggest yes (it looks continuous – same accidents), but the doctrine says the underlying reality is utterly changed (new substance). Our system can accommodate this subtlety by observing: the *provenance tag* of the bread remains the same through the change (it's a continuous physical process, no break in the thing's existence in one sense), **but** the *Constitutive nature set* changed (from {Bread} to {Christ's Body}) at consecration. Thus $ \mathrm{Con}$ before and after are different, violating the identity condition. Therefore

`Id(before, after)` is **not licensed** – the bread is not strictly identical to the consecrated host, since one had nature "bread" and the other "Christ's body." However, we can say a lot about their relation: the accidents (which are tied to appearances, not part of Con set) continue, so there is an *integrity* or continuity of accidents bridging them (this could be formalized as a secondary integrity notion). Thus, an observer perceives continuity, but by our logic *appearance-continuity alone isn't sufficient for identity* (integrity is defined in terms of the substance's persistence, not mere accidents, in this context). So we correctly predict: the senses can't guarantee the identity here because the actual Con profile changed. This provides a logical explanation of why the consecrated host is not "the same bread" even though all observable qualities seem the same – an issue medieval theologians labored over. In our system, it's simply that identity's truth-makers were disrupted: same tag, yes (it's the continuation of that piece of matter's history), but not same nature set, so identity fails. The host is a new kind of entity (Christ's substantial presence). Yet, the system can still say the host is "constituted under the appearances of bread" to explain why it looks the same – a kind of $◇$ relationship holding at the accidental level. In short, transubstantiation can be described as *change of Con with continuity of accidents*, which our logic can express (change in Con blocks identity; continuity of accidents explains why empirically one might mistake it as same [44]). Again, no special theological exception needed; it falls out from general principles of how identity is licensed or not.

These payoffs illustrate that LFI-π, with the PEI-L surface, is **general** and **uniform**. We do not solve puzzles by one-off tricks but by the same strategy: refine the conditions under which we say "is the same." It is remarkable that the **very same formal tools** (tags for origin, typed predication, guarded substitution) address puzzles in secular metaphysics and in theology [45] [46]. This suggests we have isolated something fundamental about how we track identity: essentially, we do so through historical provenance and continuity of being, whether we're talking about ships or souls. By internalizing those considerations into the logic, we ensure **consistency across cases**. Indeed, the framework offers more than piecemeal solutions – it offers a *unified account* connecting what might seem disparate issues (the one and the many, change and unity, person and nature).

Finally, an important meta-theoretical payoff: **Conservativity.** We will show (Appendix A) that any valid inference in pure first-order logic with = remains valid in our extended system, and any equation that cannot be proven in classical logic also cannot be proven in our system unless the new conditions are met. We lose no truths of standard extensional reasoning; we only add new distinctions. Technically, Id acts as an equivalence relation and congruence where it should (in purely extensional formula contexts) [21] [22]. Thus mathematicians, for example, can use this logic without even noticing a difference – it behaves when restricted exactly like the old logic (identity of numbers is just tag equality of numbers, which is built-in since numbers can be taken as having unique origins by definition, etc.). The non-classical behavior only appears when one mixes categories (like trying to identify a person with a nature, which classical logic might allow if one treated them as objects, but here it's ill-sorted), or when one introduces scenarios of change or intensionality that classical logic alone can't resolve. In those cases, our logic *departs from* classical reasoning – and rightly so, to avoid paradox. But where classical reasoning is sound, we can recover it entirely. This conservativity ensures that adopting LFI-π is not a hostile takeover of classical logic but a conservative extension – you can reduce back to classical equality talk in all the safe situations, essentially by forgetting the tags and types.

To sum up the claim: **Licensed identity (LFI-π)**, with its surface language **PEI-L**, is a paradigm shift that moves identity from the category of "given, atomic truth" to "derived, conditional truth." This shift resolves puzzles of coincidence, persistence, opacity, and doctrinal consistency not by multiple ad hoc fixes, but by a single logically coherent method. The classical equivalence relation $=$ is not abolished but becomes the shadow of a deeper relation that knows about origin and integrity – where that deeper relation is total (e.g. in a set-theoretic universe of discourse), the shadow is equivalent to it and all usual laws apply [47] [24]. Where the deeper relation is partial or qualified (e.g. cross-world or cross-category or cross-time comparisons), the system says: "no identity, but perhaps a weaker relationship applies." This not only solves individual puzzles but also demonstrates something philosophically significant: **numerical unity ("being one thing") is not a primitive fact but an outcome of prior unities (of source, of form, of continuity)** [2] [29]. That thesis aligns with ancient metaphysical intuitions (e.g. Aristotelian and Thomistic ideas that *what makes a thing one* is its form and act of being, ultimately given by God [48] [49]) while honoring modern logical rigor. The rest of the paper fleshes out this system: its background (Section 2), formal core (Section 3), surface syntax (Section 4), semantics and metaphysical options (Section 5), case studies (Section 6), theological applications (Section 7), a selection of objections answered (Section 8), and placement among related work (Section 9), before concluding (Section 10). Appendices provide technical backup (proofs of conservativity, axioms, models, and a glossary and guide for readers). We proceed now to background and motivations, to firmly set the stage for the technical development.

2. Background and Stakes

2.1 Classical identity in brief: Axioms, laws, and where it shines

The classical concept of identity in logic is encapsulated by a few simple axioms. In any standard first-order logic with equality, one assumes: - **Reflexivity:** $\forall x,\; x = x$. Everything is identical to itself. - **Symmetry:** $x=y \to y=x$. (Often derivable from substitutivity, but conceptually part of identity being an equivalence relation.) - **Transitivity:** $x=y \wedge y=z \to x=z$. Together these mean "=" is an **equivalence relation** partitioning the domain into self-same units. In addition, one typically includes an axiom schema for **Leibniz's Law**, also known as the **Indiscernibility of Identicals**: $$x = y \to (F(x) \leftrightarrow F(y)),$$ for any property or formula $F(-)$ in the language [5]. In practice, this is used in the contrapositive form: if there is *any* predicate or property that differentiates x and y ($F(x)$ true and $F(y)$ false), then $x \neq y$. Equivalently, if $x=y$, then they share all the same properties. This is a very powerful principle, essentially ensuring that there cannot be two distinct entities with all the same properties (sometimes called the **Identity of Indiscernibles**, which is a closely related principle often attributed to Leibniz as well).

There is also the **Substitutivity of Identicals**: if $x=y$, then one can replace x with y (or vice versa) in any **extensional** context salva veritate (preserving truth). In formal proofs, one uses a rule: from $x=y$ and any statement $\phi(x)$, infer $\phi(y)$ (provided ϕ is a formula where substituting y for x makes sense and does not change the meaning). This rule is basically a direct consequence of Leibniz's Law if we treat $\phi(-)$ as a property that is true of x; then $x=y$ implies $\phi(y)$ must hold if $\phi(x)$ did. In sum, classical identity means **absolute indiscernibility and unrestricted substitutability**. These features ensure that identity in classical logic is a **very restrictive relation**: it can only relate a thing to itself. As David Lewis quipped, under the classical view it seems each thing trivially stands alone: "each thing is identical to itself and to nothing else" [3]. There is no nuance—either two names refer to one object (if all properties match), or they don't (if any property differs).

In everyday and scientific practice, this classical notion of identity works exceedingly well in **extensional contexts**: - In mathematics, we rely on a strict notion of equality: if two numbers are equal, any equation or statement about one holds for the other, etc. We don't encounter "slightly equal" or "equal in some respects" – equality is clear-cut, and mathematics is largely extensional (no hidden context where substituting 7 for

3+4 causes problems, since $3+4=7$ absolutely). - In set theory and most of formal semantics, identity is the backbone of interpreting statements like “there is at most one x such that …” or “x is the same as y”. Without classical identity, we couldn’t even formulate uniqueness conditions properly. The rigidity of identity ($=$) provides a solid foundation for these. - Leibniz’s Law is intuitively obvious in many cases: If I say “Mark Twain is Samuel Clemens” (which is true, since they are one person), it follows that anything true of Mark Twain (being an author of *Tom Sawyer*) is true of Samuel Clemens, and vice versa. Failing to accept that leads to absurdity if we’re talking about the same individual. Classical identity shines in such cases: it forces us to treat them as one entity in inferences, as we should. - The classical theory easily explains certain falsehoods: e.g., “The morning star = the evening star” turned out to be true (both are Venus), so before that was known, people were effectively treating them as distinct and could find a property (seen in morning vs. seen in evening) to differentiate – hence not identical by Leibniz’s Law until further knowledge collapsed those properties. Once it was discovered they are the same planet, the apparent property difference was explained away (one property was being observed at different times). Classical identity wasn’t *wrong* there; it was our information about properties that changed. - Identity also undergirds the notion of **counting and individuation**: if I say “there are 3 apples on the table,” that presumes I have a way to identify which things count as the same apple or a different one. Classical logic assumes that identity is a given equivalence relation in the domain, so semantics can treat “3 apples” as meaning there exist at least three distinct (non-identical pairwise) entities that satisfy “is apple on table.” Without a clear identity relation, even the semantics of numbers and quantifiers would be murky.

In short, classical identity provides **clarity, consistency, and a simple framework** whenever we have a well-defined domain of objects that do not change in problematic ways or appear under different guises that matter. It “shines” especially in: - **Extensional mathematics and logic:** where identity of individuals is unambiguous and all contexts are transparent (no hidden modalities or propositional attitude contexts). - **Administrative and legal contexts:** where identity is often treated in a straightforward way (if two records have the same social security number and matching details, they identify the same person; if not, they are different people – though even law can get tricky with identity in cases of e.g. identity fraud or change of legal identity). - **Common sense tracking:** Usually, if I see my friend today and see someone looking and acting exactly like them tomorrow, I take it for granted it’s the same person. Classical identity fits this everyday reasoning until we hit something like identical twins or mistaken identities, which are then resolved by noticing some differing property.

Philosophers from **Leibniz** onward have emphasized these formal virtues. Leibniz himself formulated the classic laws and pointed out that if an identity statement is true, then it is necessarily true (assuming the entities are *really* identical, you cannot have had them be distinct) [50] [51]. This is sometimes called the **necessity of identity** (if $x=y$, then in every possible world $x=y$). Again, in classical logic this holds by semantic reasons (if they’re the same object, they’re the same in all worlds under standard semantics for constants or for rigid designators).

So classical identity gives us a **robust, all-or-nothing sameness** that fits many purposes. The problem is not with its internal consistency or usefulness in its domain – it’s with the *limits* of that domain. When identity gets entangled with *modal reasoning* (possible worlds) or *intensional contexts* (belief, knowledge), or with *trans-world or cross-time comparisons*, things become trickier than the pure logic can handle without supplementation. But before addressing those troubles (Section 2.2), it’s worth noting that historically, many philosophers thought those additional contexts required more nuance.

For example, Aristotle distinguished **numerical identity** from mere **qualitative similarity** very clearly [52]. He grappled with the question of what makes a thing one and the same through change (the *unity of a substance*). Medieval philosophers like **Avicenna** and **Aquinas** introduced notions like the "individuating principle" (e.g. *materia signata* – designated matter – in Aquinas) for material substances, essentially trying to give an account of identity's ground beyond just form [53]. They didn't have formal logic symbols for identity, but they were dealing with the puzzle of how you can say "this is the same thing as it was before" when obvious changes happened. Already they intuited that something like *continuity of existence* or *same underlying matter/form* was needed as an explanation.

Leibniz's own **Identity of Indiscernibles** (the idea that if two things share all properties, they must be one) is not derivable in standard first-order logic (it's consistent to imagine two distinct objects that just happen to have all the same properties, unless you have a strong background principle ruling that out). But often a form of it is assumed metaphysically: e.g. if you truly cannot find any difference between A and B even in principle, why treat them as two? Our approach in LFI-π resonates with that by saying: if all the relevant grounding features line up (origin, nature, continuity), then indeed they collapse into one via `Id`. If any differ, they remain two. Leibniz also insisted on the **necessary truth of identity**: if A = B, then necessarily A = B [50]. In our system, something similar will hold within the extensional fragment; but we'll also see that one can speak of identity relative to a particular causal context (so cross-world identity becomes subtler – something our framework can handle by parameterizing provenance by world, for instance).

In conclusion, classical identity is a simple and powerful tool – until it isn't. It shines in homogeneous contexts with no ambiguity about criteria of sameness. As the famous line goes, *"for any x and y, either x is y or x is not y"* (the law of excluded middle applied to identity) holds tautologically, but the real philosophical action is in *deciding* or *explaining* the cases where we initially aren't sure if x = y. That's where classical theory just says "check all properties" or "it's a primitive fact" and leaves puzzles unresolved. The next section reviews why those puzzles bite hard.

2.2 Why the puzzles bite: Constitution, persistence, intensionality, mixed predication

We have already previewed the major puzzles in the Introduction, but let's analyze a bit deeper *why* they pose a challenge to the classical identity framework:

- **Material Constitution (Statues, Clay, etc.):** At stake is the difference between an object and the material that makes it up. Classical logic has no built-in distinction between an object and its matter; if at time t_0 the statue and the lump coincide entirely (same spatial region, same matter particles), then one might think they are identical. Yet they appear to have different properties (**modal properties** especially: the clay could survive being reshaped, the statue could not [54] [55]; or temporal properties: the clay existed before the statue was formed, if it did, etc.). These violate Leibniz's Law, so classical logic would say they are not identical. Fine, so they are distinct – but then how can two distinct things be in the exact same place at the same time, made of the exact same matter? That offends our intuitions about physical objects (it seems like double-counting what is really one thing). This is the **coincident objects** problem [54] [55]. Some philosophers (e.g. David Wiggins, Judith Thomson, etc.) have argued that we have to carefully differentiate the *sortal concepts*: "statue" vs "piece of clay". They might adopt a form of **Relative Identity** ("the statue and the clay are the same material object but not the same artifact," etc.). Relative identity theories (Geach's idea) basically say identity is always relative to a category. But classical logic cannot formalize a relative identity easily without giving up on "=" as absolute. Our approach in effect provides a way to say "same in some respect (matter) but not absolutely same" using $\Diamond$ (constitution) for the material aspect sameness and withholding Id. The reason the puzzle *bites* classical identity is because classical identity is too crude to express "sameness in one way but not another" – it's all or nothing. Thus it forces a choice that both options of which seem wrong: either the statue = the clay (then how do they differ? one must deny the difference or some property), or statue $\neq$ clay (then how do they coincide? one has to allow co-location of distinct objects). There's no middle ground in the classical scheme, whereas intuitively, the truth is in the middle: they share something deeply (the matter) but not everything (form or essence). This is exactly the kind of nuance our "licensed identity" tackles by separating conditions (same origin, same form, etc.). Classical identity provided no direct way to handle that beyond metaphysical arguments outside the logic (like "maybe form and matter are different aspects").

- **Persistence and Change (Ship of Theseus, Personal Identity):** Here the issue is identity *over time*. Classical logic itself is atemporal – $x=y$ doesn't say at what time, it just says two designators pick out one thing. To reason about identity over time, one typically has to either incorporate time indices (saying $x(t_1) = y(t_2)$ meaning the entity x at time1 is the same as entity y at time2), or use modal logic or second-order logic to talk about trajectories. The Ship of Theseus scenario reveals a problem: classical logic can represent the ship at time0, the ship at time1, etc., and ask "are they the same?" But if at time2 we have two claimants (the continuous repaired ship vs. the rebuilt original-parts ship), classical logic alone can't decide because purely qualitatively each has some claim: one has continuity of form and ownership, the other has continuity of original parts. As soon as you try to apply Leibniz's Law: they are obviously discernible (one has mostly new planks, one has mostly old planks), so logically they are not identical. But which one is identical to the original? If we say the original = continuous one, then original $\neq$ reassembled one; if someone else says original's "true identity" followed the parts, then original = reassembled, $\neq$ repaired. There's no logical contradiction here because those are different criteria being chosen. The classical framework doesn't tell us *which* criterion to use – it just says whatever criterion is correct must be such that it preserves all properties (so pick the right property set). The puzzle bites because we have a **branching** – two candidates for one predecessor – and classical identity disallows one thing being identical to two. So at best, one of them can equal the original. But nothing in classical logic favors one choice over the other; that depends on how you weigh continuity of form vs. continuity of matter. This becomes somewhat subjective or at least extra-logical. The classical theory has to remain silent or push the criterion into the meaning of "same ship" implicitly. Our framework bites the bullet by building in an **anti-branching rule**: essentially it chooses by fiat that only one line can carry the original identity. Which one? That depends on how we formalize integrity: e.g. if we say integrity is mainly about form/function continuity (not caring about original atoms), then the continuously used ship wins (it keeps the tag π and the other gets a new tag). If we insisted integrity is strict material continuity (no tolerance for replacement), then perhaps by the time enough planks are replaced, one might argue the ship lost identity until maybe the reassembler gives it back – but that gets complicated. The key is, classical logic couldn't even talk about "no branching" because it lacked the notion of an *identity trajectory*. We introduce that via tags and NB. The puzzle also bites in cases like **personal identity fission** (Parfit's teleportation or split brains): classical identity can't allow one person to become two, so it says either the original died or only one of the two is "really" him. Those conclusions feel inadequate (who's to say which one is him if both have equal claim?). Many have argued for more relativized notions (like "half-identity" or just denying that identity is what matters for survival, as Parfit famously did). Our system's take is: indeed, you can't have one become two under strict identity, but you can have a period of indeterminacy or rather a *branching equivalence* that then ceases to be identity. One can formalize Parfit's notion that identity isn't what matters by noting that in fission, Id stops applying but a weaker relation ($≡_π$ perhaps without NB) still links the persons, which might correspond to psychological continuity. The classical scheme had no place for a "halfway" relation except to say "they're not identical, but something else is true." We make that "something else" explicit (common provenance, etc.).

- **Intensional Contexts:** These include modal contexts ("necessarily, …"), epistemic contexts ("Alice believes …"), etc., where substitution of identicals can fail. The classic example given earlier: Clark Kent vs. Superman in Lois's belief. Why does classical identity struggle? Because $Believes(Lois, \text{Superman~can~fly})$ is not a truth-functional context; Lois's belief depends on how she conceives the person, not just the referent. In classical logic, one might represent belief via a relation $B(Lois, x\text{-flies})$ or so, but to capture the failure of substitution you typically need a more advanced logic (like modal operators for belief, which don't allow substituting co-referential names not known to the believer). The **failure of substitutivity** in such contexts shows that Leibniz's Law has a ceteris paribus: it only holds when F is not an intensional context that can make a distinction between x and y being presented. In practical logical systems, one restricts Leibniz's Law to **extensional properties** only. But standard first-order logic doesn't have a built-in notion of intensional vs extensional positions – everything is extensional by assumption. Handling belief or modality often means going to a modal logic where $=$ might not behave the same inside a modal operator (you might treat "=" as necessary equality or not). The puzzles bite because we have true premises like "Clark = Superman" and "Lois believes Superman can fly" yet a false conclusion "Lois believes Clark can fly". Classical first-order logic with identity would allow that substitution and derive a falsehood, so one has to prevent that by weakening the logic or adding ad hoc rules (like an opaque context rule). Our approach solves this by a simple idea: **belief contexts are not extensional positions** for substitution unless the thing believed is extensional. We effectively will not license the substitution in Lois's belief context because the provenance or "evidence" conditions aren't met from her perspective. More concretely, we might index belief by conceptual tags: Lois's concept of Superman vs her concept of Clark are different "provenances" in her epistemic structure, so in the logic we cannot assert $Id_{\text{Lois}}(Clark, Superman)$ in her belief context (no common epistemic tag). We don't delve deeply into epistemic logic here, but the framework is adaptable to that. For modal necessity contexts, consider "necessarily $a=a$" (true) vs "necessarily $a=b$" (not true even if actually $a=b$, unless one subscribes to the necessity of identity as an absolute). In modal logics, if a and b are rigid designators and $a=b$ in the actual world, then necessarily $a=b$ is usually taken as true (Kripke's point) [50]. However, some intensional contexts like counterpossibles or hyperintensional distinctions might still differentiate. Our logic's treatment of modality would likely involve tagging across possible worlds, which is beyond scope, but the main takeaway: intensional contexts require *guarding substitution*, which our licensing idea provides by saying: we only substitute when a certain relation holds (here, probably something like "same referent known under the same mode of presentation"). Classical logic does not natively handle modes of presentation – we introduce tags that can encode such modes.

- **Mixed Predication (category mistakes and theological "is"):** This refers to statements where the subject and predicate seem to be of different categories, or where "is" doesn't mean simple identity but something else. In theology, *"Christ is God"* is a classic case of *mixing a person and a nature* in one statement. If taken as identity ($Christ = God$?), it would imply a person (supposit) equals the divine nature itself, which in orthodox doctrine is not correct – the Second Person of the Trinity is not *identical* to the divine nature, He *has* the divine nature (and so do the Father and Spirit). Similarly, *"Christ is human"* is not saying Christ = humanity itself, but that He possesses human nature. These are predications of a nature to a person. Classical logic with one sort for all individuals can't differentiate whether "is" denotes set membership, identity, predication, etc., aside from context. It would typically interpret a sentence "X is Y" as either $Y(X)$ (Y as a predicate) or $X=Y$ if Y seems like a name. The danger is if one treats "God" as a name of an entity and "man" as a name of an entity, then "Jesus is God" looks like an identity claim between two entities – which yields the Trinity problem of then Father = Son, etc. If one treats "God" as a predicate (meaning "is divine"), then "Jesus is God" can be formalized as $Divine(Jesus)$, and "Jesus is man" as $Human(Jesus)$, which is fine. But then how to formalize "Father, Son, Holy Spirit are one God" is trickier—if "God" is not a person but a nature, one can say they share one nature (like $Divine(F), Divine(S), Divine(Sp)$ and uniqueness of divine nature), but classical logic doesn't have a direct way to enforce "one nature in three persons" except via sets or higher-order. In any case, these issues bite because of **category differences**: person vs nature, whole vs part, etc., where using identity would collapse distinctions. Our multi-sorted approach with S (supposits/persons) and N (natures) sorts already avoids certain mix-ups: a well-typed equation can only be between two S or two N entities, never between an S and an N. So one cannot even write $Jesus = DivineNature$ if those are different sorts – it would be a type error. This is how the logic enforces "no person-nature identities" by design [26]. Mixed predication like "Jesus is God" is parsed as a supposit-nature predication $Jesus \in Divine$, not an identity. Classical logic didn't have that tool without adding ad hoc predicates or sort distinctions. Thus the puzzle of communicatio idiomatum (how can we say "God was born" meaning a divine Person in human nature was born, without meaning the divine nature was born) bites classical logic because it lacks an aspect mechanism. Our system directly addresses that with typed predication and the aspect guard rules (only allow certain substitutions if aspect matches).

Overall, the puzzles bite classical identity theory because **classical identity has no structure**: it's a two-place relation with only the trivial structure of being or not being. It doesn't account for time, for modality, for aspect, or for any internal composition of identity. As a result, whenever an identity question depends on *other factors* (like origin, continuity, category), classical theory must either ignore those factors (leading to paradox) or say "identity doesn't hold" (which sometimes seems to conflict with other intuitions of unity). The stakes are high: if we keep classical identity as is, we tend either to live with the paradoxes or to create piecemeal modifications (relative identity, sortal-dependent identity, temporal parts theory, counterpart theory, etc.). Each of those solutions adjusts something: relative identity says there are many "=" relations (each sortal has its own); temporal parts theory says an object at one time is only part of a four-dimensional worm so identity over time is being part of the same worm; counterpart theory (Lewis) says "is the same in another world" is not literally identity but a counterpart relation. Our approach could be seen as one more contender: **licensed identity**. But an advantage is that it tries to capture the insights of several of those within a single formal system: like sortal relativity (via Con sets and I conditions), temporal continuity (via I and NB), origin essentialism (via R tags), etc., all integrated. The puzzle cases will be revisited in Section 6 with plain-language explanations of how the calculus deals with them.

2.3 Participation metaphysics in one page: cause of unity, form reception, integrity across change

Behind our formalism lies a certain philosophical worldview inherited from classical metaphysics, especially the notion of **participation**. Since we call our approach "participation-friendly," we should briefly sketch what that means. In classical philosophy (Plato, Aristotle, and particularly the Neo-Platonic and Scholastic traditions), to say a creature *participates* in something means it **receives** some share of a perfection or nature from another, more fundamental source. Applied to the question of identity and unity:

- **Cause of Unity:** Traditional metaphysics asks *what makes a substance one?* For Aristotle, a key part of the answer was the **substantial form** – the organizing principle that makes matter into a single entity of a certain kind. For example, my body is a unified organism because the form of human life structures it. Aquinas builds on this: a thing is one because it has one act of existence (which it receives from God) and one substantial form informing its matter [56] [49]. In this sense, the unity of a being is *caused* – it's not a given that a pile of parts is one being until a form makes it so. Aquinas even says *being* (existence) and *oneness* are convertible – everything that exists is, in some respect, indivisible (one) as long as it is a being. So the cause of the unity of an entity lies in its origin: both the **originating efficient cause** (the maker or generator) and the **formal cause** (the nature it receives). Our framework mirrors this: we say two things can be identical only if they have **the same origin (efficient cause)** and the same **form/nature**. If they differ in either, they are two. This is essentially implementing a participation idea: the origin (provenance tag π) is like the mark of having been generated by the same act or source, and the nature set Con(s) ensures they share the same form(s). So the "cause of unity" in metaphysical terms – *God's act of existence flowing into a creature's form* – is represented logically by these conditions: a common originating source and a common essence.

- **Reception of form (participation in nature):** Creatures in the Scholastic view do not have existence of themselves; they have it by participating in God's being (since God is Being itself, creatures have being by attribution) [57] [58]. Likewise, creatures are not universal essences; they are particular by matter and by having a share of the nature. For example, Socrates *participates* in humanity – he is not humanity itself (which is an abstract form), but he has humanity in him. This is the idea of *participation in a perfection*: Socrates participates in the perfection of human nature; all humans do, which makes them similar, but each is a distinct instance. Now, when we talk about identity, if two purported entities were truly one, they would in some sense be sharing *the exact same instance of nature and existence*. If Socrates and Plato were somehow numerically identical (they are not, obviously), they would literally be one instance of humanity, one soul, etc. That doesn't happen naturally – each person gets their own participation in the form. Our logic enforces that: $ \mathrm{Con}(x)=\mathrm{Con}(y)$ means they have the same nature *multiset*, but does it mean the same *instance*? Actually, if x and y are distinct supposit sorts, even if Con sets match as sets, if they are not Id, those natures are still distinct instances. Only when $Id(x,y)$ holds can we say they share the *instance*. In theology, Jesus Christ is one person with a divine and human nature – that's a unique case where one person participates fully in two distinct natures. Generally, identity means one *supposit* underlies. So participation metaphysics would say: creatures are distinct because each one has its own participation in being (Aquinas: *all beings apart from God are beings by participation* [57]). Our tags align with that: everything except God has an origin tag given by some cause (God or secondary causes) – there's no "two things with one tag" unless they are actually one and the same intended effect of that cause. We explicitly disallow two distinct individuals from having the same exact root tag (except transiently in a fission before we issue a new tag to one). This resonates with the idea of no two things have the exact same *esse* (act of being) – if they did, they'd collapse into one. So *tags* can be thought of as representing something like a bestowed act of existence or a lineage of being.

- **Integrity across change:** In participation terms, everything except God is mutable and can lose or gain properties. What keeps it the "same thing" through change? The classical answer: as long as the substantial form and primary being remain, and only accidental changes happen, the substance remains the same. If a substantial change happens (like death, which Aquinas would call corruption of the substance and loss of the soul form), the thing ceases to be itself. Thus integrity is tied to *staying within the limits of changes that do not destroy the essential form*. Different kinds of things have different tolerance: an organism can heal wounds (accidental change) but not survive decapitation (probably a substantial change for a human). A machine can have parts replaced and still be "the same machine" if it retains structure, etc. Our notion of *integrity policies* captures that idea: it's essentially formalizing the threshold between mere change and destruction of identity. For any sort of entity, we specify what kind of changes it can undergo while still "participating in the same act of being" (keeping the same tag). If changes exceed that (like branching into two, or losing the essential form), then the original participation ends and a new one begins (new tag). For instance, if a cat dies, we'd say the corpse is not the same entity as the living cat – in our system, perhaps at death the tag terminates (the life principle is gone), and the matter is now under a different kind of form (just cells decaying). So we wouldn't have $Id(\text{cat before}, \text{corpse after})$. Instead the corpse might get a new tag as a separate object (if we even choose to count it as an "object" rather than just remains). Another example: in Sacramental theology, bread's substance changes to Christ's flesh; that's a substantial change, so identity is lost because Con changed (like we described). So integrity is about *continuity of participation*: does this thing continue to have the same share in being that it had? For material things, often continuity in time and space (no gaps, no duplication) plus same form suffices. For persons, psychological or bodily continuity might be needed on top of physical. Our logic doesn't decide which; it just provides the slot $I(x,y)$ to be filled by the appropriate condition per category.

A final point: **God does not participate.** In classical doctrine, God is *ipsum esse subsistens* – existence itself, not a participant in a greater being [57] [58]. All other things have being by sharing in God's creative act. In our framework, we reflect this by saying: God (the Divine Nature, or the Persons in their divine existence) does not come with an origin tag given by another, because God has no cause. Thus, when we model the Trinity, we gave the Father tag "U" to indicate unoriginate (not from another). The Son's tag "F" meaning from Father, but note: in a sense, that's an eternal procession, not a "creation tag" – it behaves differently (the Son's tag might better be thought of as an eternal relation of origin, not a temporal cause). We will have to treat divine tags specially – they don't imply a higher cause, just relational origin. Likewise, integrity for God is trivial; God cannot change or lose the divine nature. So in the theological model, we don't apply the same integrity rules to God (God is immutable and simple). Essentially, the divine nature stands outside the participation framework; creatures participate in God's perfections analogically [59], but God just is those perfections. Our system acknowledges this by never treating the Divine Nature as a sort that needs a tag (we don't say the Divine Nature has a cause or tag; the Persons have relational tags just to differentiate them). In technical terms, our model of the Trinity will not bind the Persons under a common provenance that would force identification – each Person is marked as distinct by origin precisely to preserve that they are really distinct, while all three participate *univocally* in the one Divine Essence (meaning literally the same essence, not just a similar one) [60] [59]. Creatures, on the other hand, can only participate in God's attributes analogically (they can be wise or good, but not Wisdom or Goodness itself). The logic doesn't explicitly enforce analogical vs univocal participation, but we mention it here conceptually because it influenced our thinking in adding a participation modality in Section 5.3. We will allow a predicate $Part(s, F, m)$ to say "s participates in perfection F in mode m (univocal or analogical)". This is not in the core logic LFI-π, but an add-on that ensures, for example, no creature is univocally participating in the Divine Nature (that would make it God); only the Persons share it univocally among themselves [59].

To summarize this section: The metaphysical backdrop is that *unity comes from shared causes and forms*, not from a bare logical predicate. By aligning our formal criteria with metaphysical ones (same cause, same form, continued existence = one being), we ensure that our "licensed identity" is not an arbitrary formal tweak but is capturing an intuitive order of explanation: **existential continuity precedes identity** [29] [61]. As one might say, only God is absolutely one by Himself; creatures achieve unity through participation in existence and form, and thus two creatures cannot be one unless they literally share those – which in effect means they are the same creature. We reflect that by requiring those shares to match for identity. This provides an ontological rationale for the logical rules: we're modeling logically what many philosophers implicitly or explicitly pointed to when discussing unity (though they didn't formalize it like this). It's a "participation-friendly" logic because it respects the difference between *having* a nature and *being* a nature, and between receiving being and being Being.

(*As an aside for readers: you don't need to buy into the entire metaphysical picture to use the logic – the logic can be viewed simply as a formal system with certain extra relations. But it's helpful to know that it wasn't concocted arbitrarily; it stands on the shoulders of a long tradition grappling with* the one and the many*. This motivates our design choices – such as multiple sorts for persons vs natures, the use of tags (a bit like principium individuationis), and the emphasis on non-branching continuity.*)

2.4 Doctrinal constraints for easy reading: One supposit of Christ, distinct Persons by origin, one essence and will, inseparable operations

Because we will apply our logic to theological doctrines (Trinity and Incarnation in particular), it's useful to lay out in plain terms the orthodox constraints we must respect. This will serve as a checklist to show that our formal system indeed does justice to these, and it will help readers unfamiliar with the doctrines to follow those examples without getting lost in technicality. The key doctrinal points are:

- **Christology (Incarnation):** In Christian doctrine, Jesus Christ is one **person** (one *suppositum*, in our terms) who has two complete **natures**: divine and human. This is defined classically by the Council of Chalcedon (451 AD), which taught that Christ is to be "acknowledged in two natures *unconfusedly, unchangeably, indivisibly, inseparably*… the distinction of natures being in no way annulled by the union, but rather the properties of each nature being preserved, concurring in one Person and one hypostasis" [11] [62]. In other words, the divine nature and human nature remain distinct in Christ (he doesn't morph them into a single mixed nature), yet they belong to the one Person of the Word. There is no *human person* separate from the Divine Person; the Word assumes impersonal human nature. The upshot: *there is exactly one supposit (the Word) and it has two natures.* Any logic we use should be able to talk about that without either splitting Christ into two or mixing the natures into one. Additionally, by the **communication of idioms**, we can say things like "God was born of Mary" or "Jesus (as God) is eternal" or "Jesus (as man) hungers" and those be true in the right sense. Formally, it means the logic should allow that the *subject* (the person) is one, while predicates can be attributed according to either nature, with the understanding of which nature they come from [10]. We must avoid a situation where "Jesus = God" and "Jesus = a man" lead to "God = a man" by transitivity – our typed system will prevent that by not equating person with nature. Instead, "Jesus is God" will be personal predication of the divine nature, not an identity. Also, Christologically, it's important that Christ has **two wills** (divine will and human will) corresponding to his two natures (this was clarified in the 7th century against Monothelite heresy). Our logic doesn't explicitly model wills, but if we needed to, we could treat will as a property of nature so that Christ can have a human will and a divine will without conflict. The phrase "one person, two natures" is the guiding light, and logically we adhere to that by making sure in any representation of Christ, we have one S with two N in its Con set.

- **Trinity:** Orthodox Trinitarian doctrine says there is *one God in three Persons*: Father, Son, Holy Spirit. They are distinct *Persons* precisely by their relations of origin (Father unbegotten, Son begotten of the Father, Spirit proceeding from Father (and Son)) [39] . They are not three gods because they share one divine *essence or substance*. "Consubstantial Trinity" is the term: each Person *is* fully God (has the divine nature entirely), not a part of God [36] [37] . They also have one divine **will** and one **power** and hence perform all external actions (creation, salvation, etc.) unitedly (this is the doctrine of inseparable operations) [40] . In the words of an early creed: *"He is not the Father who is the Son, nor is the Holy Spirit he who is the Father or the Son... They are distinct from one another in their relations of origin... The Father generates, the Son is begotten, the Holy Spirit proceeds... but the Father, Son, Holy Spirit are one in omnipotence, majesty, and substance."* [39] [63] . For our logic, that means: we need to represent three distinct S (supposits) that all have *Divine* in their Con set. We must ensure that no two of them get accidentally identified – which we do by giving them distinct provenance tags (there is no single π that both Father and Son share exclusively, since one is unbegotten, the other begotten). We also represent that each one $\in$ Divine and perhaps assert a principle that there is exactly one divine nature instance (if one were formalizing, one could say something like any two divine supposits have Con set that includes the same *numerical* nature – effectively they share an essence instance, though how to formalize that is subtle; often one can treat "Divine nature" as a single abstract object and say all three constitute it). The logic should natively allow "Father is not Son" etc., which it does by sorted identity: they are distinct S with different origin tags, so Id never identifies them. Meanwhile, a statement like "Father, Son, Holy Spirit are one God" can be interpreted as "there is one divine nature that Father, Son, Spirit each are/constitute". We might formalize "one God" as "there is exactly one x such that x is a Divine nature instantiated in these three persons". Our system certainly ensures one *kind* of oneness: if we quantify over Persons, there are three; if we quantify over Natures, there is one Divine nature kind. The **one will** part implies that any predicate like "willing X" which is proper to the divine nature (an operation of God) should not differentiate the Persons: if the Son wills something, so do the Father and Spirit, because there's numerically one volitional power. We can enforce that by treating "wills X" as a predicate true of the divine nature, and then by our distribution law, it will be true of any supposit constituting that nature [40] . And inseparable operations in general: if God (the divine nature) does some external act, then Father, Son, Spirit all do it together. Our logic will allow that because we don't require distinguishing which person does an act that's at the nature level. Internally, the persons do have distinct roles in internal (immanent) actions (generation and procession), but those aren't "ad extra" operations, they are relations of origin, which we mark via the tags.

- **Other doctrinal notes:** The system should also allow *"distinct persons, one essence"* without conflating person and essence. It should allow statements like *"The Father is God, the Son is God, the Holy Spirit is God"* to be true (each person has the divine nature) but *"The Father is not the Son"* also true [39] . It should allow *communicatio idiomatum* statements about Christ as we discussed. It should also allow or at least not contradict other theological principles like: *"the divine persons are only distinguished by their relations of origin"*, *"the whole Trinity is involved in every external work"*, *"the Son of God died in the flesh"*, etc., which we can parse appropriately. Our logic's discipline basically ensures no category mistakes: we never say a divine person = the divine nature or a human nature = the divine nature, etc., thus avoiding the heresies of Sabellianism (confounding persons) or Eutychianism (confounding natures).

By laying these out plainly: One supposit Christ with two natures; Three divine persons with one nature, one will, one action externally; Persons distinct by origin; No mixing nature and person in identities – we can be confident our framework will be *orthodox-compliant*. Indeed, one of the motivations of this work was to **regiment theological language** in a logically coherent way so that, for example, we do not inadvertently prove a contradiction from true premises of faith (like proving "God is not God" by bad substitution). The system effectively builds those centuries of careful distinctions into the logic rules so that such moves are blocked by design, not just by informal caution.

With this background and these constraints in mind, we now turn to the formal development of the core theory, LFI-π. In the next section, we will define the formal language, the new predicates (like $R_π, ∈, ◇, Id$), and the axioms/rules that govern them. We'll see explicitly how identity is defined from provenance and integrity, and what inference rules change in comparison to classical logic. After that, we will revisit how the surface notation (PEI-L) makes it easier to work with, and then explore semantics and applications.

3. Core Theory: LFI-π (Licensed Formal Identity with provenance tag π)

3.1 Language and typing

We begin by specifying the formal language for the core calculus, LFI-π. This is a **many-sorted first-order language** – meaning we have multiple domains of discourse with different sort identifiers, to prevent category errors. Specifically, we employ the following sorts: - **S (Suppposits):** The sort of individual subjects (also called *suppositum* in theology, roughly meaning a concrete individual substance or person). Elements of sort S are things like particular people, animals, material objects – anything that can serve as a bearer of natures or properties. In our theological context, divine Persons are also elements of S (they are supposits, albeit immaterial ones). We use lowercase letters like s, t, a, b (possibly with subscripts) as variables for sort S. - **N (Natures):** The sort of natures or substantial forms. These are like types or essences that a supposit might have. For example, *Human*, *Divine*, *Feline*, *Statue*, *Clay* can be thought of as elements of sort N (or maybe as predicates we will reify as N-sort constants). We use lowercase n, m, N, M for variables ranging over natures. We treat natures as abstract objects in their own right in the logic (one could alternatively make a predicate for each nature, but here it's useful to have a sort so we can talk about relationships between natures, like one nature being in a constitution set equal to another's). - **$Π$ (Origin Tags):** The sort of provenance indices or tags. Think of these as identifiers for "histories" or causal chains. Each supposit is associated (at least piecewise) with some $π ∈ Π$ that marks its origin or lineage. Tags could be structured (like "Father" for Son's origin), but formally we just treat them as atomic symbols or values of sort $Π$. We use Greek letters π, ρ, σ for these. In models, you can imagine a function that gives each supposit its tag(s). There may be multiple kinds of tags (we could have tags for different causal aspects, e.g. material origin vs formal origin), but for now assume one primary origin tag per object, plus possibly special tags like "U" (unoriginated) for God. - **T (Time or indices):** We include an index sort T for temporal indices (or could be more general indices like world-state indices if doing modal stuff). This allows us to talk about things like $R_π(s,t)$ holding *during* a certain interval, etc., or to have time-varying predicates. For now, we can think of T as time moments or states. Variables like i, j, t_1, t_2 range over

T. If one is not dealing with time, this sort can be collapsed or ignored, but we include it to handle persistence formally.

With these sorts, we now list the non-logical symbols of the language: - **Predication:** A binary predicate $P(s,n)$ (or we might write $s ∈ n$ in surface syntax) of sort $S \times N$. The intended meaning is *"supposit s truly has nature n"*. For example, $P(Socrates, Human)$ means Socrates is human (possesses human nature). P corresponds to what logicians sometimes treat as an `$s$ is an $n$` atomic formula. We will often use the infix $\in$ in the prose, but formally it's just a predicate symbol. - **Constitution:** A binary predicate $K(s,n)$ (or surface $s◇n$) of sort $S \times N$. This means *"supposit s is constituted under nature n"*. The difference between P and K is subtle but important: $P(s,n)$ is like saying s has nature n in actuality (a true predication), whereas $K(s,n)$ might be seen as saying n is one of the essential natures that make up s's being (whether or not s fully acts according to that nature). In many cases P and K will coincide (if s has only one nature, then to say $s◇N$ is basically to say s is of nature N). But if an entity has multiple natures (e.g. Christ), one might say $K(Christ, Human)$ and $K(Christ, Divine)$ both hold, meaning Christ's constitution includes both humanity and divinity. Meanwhile, $P(Christ, Human)$ and $P(Christ, Divine)$ are also true – He truly is both. If a being temporarily or potentially has a nature, one might use K differently (like a soul that could have a body, etc.). But for our purposes, $K(s,n)$ primarily functions to collect the set of n's that belong to s. We define $\mathrm{Con}(s) := { n \mid K(s,n) }$ as the set of natures constituting s. In axioms, we may impose that if $K(s,n)$ then $P(s,n)$ (constitution implies actual predication), ensuring no "dormant" natures. We will indeed adopt that: *constitution strengthens predication* (so an $◇$ implies an $∈$)【4.3】. - **Provenance/Origin:** A ternary predicate $R_π(s, s')$ or $R(s, s', \pi)$ of sort $S \times S \times Π$ (with perhaps an implicit or explicit time parameter). We often write it as $R_π(s,t)$ meaning *"s and t share origin tag π"* or *"s has provenance π and so does t."* In effect, $R_π$ is like an equivalence relation on S labeled by π. We could have done a function $\tau(s) = \pi$ giving the tag of s, and then define $R_π(s,t) := [\tau(s) = \pi = \tau(t)]$. In fact, in models we often do that [19]. But in the logic we treat R as a primitive relation, with axioms that make it behave like: for each π, $R_π(-,-)$ is an equivalence relation (reflexive, symmetric, transitive) on the set of supposits that originated from π. If we quantified π, we could say $∃\pi\,R_π(s,t)$ means s and t share some origin. We will use that soon for identity definition. $R_π$ is subject to a crucial constraint: **Non-multiplicity** (no branching) which will be a separate predicate, but conceptually related: if $R_π(s, u)$ and $R_π(t, u)$ and $s \neq t$, this might indicate branching (two distinct things with the same origin). We will handle that with NB below. - **Integrity:** A binary predicate $I(s, s')$ of sort $S \times S$ (possibly also needing a time interval argument, but we can suppress it by encoding times in s, s' identity or by a separate notion of time-interval integrity). $I(s,t)$ means *"s and t satisfy the integrity/persistence conditions for being one continuous entity."* Roughly, $I(s,t)$ is true if s at time t_0 and t at time t_1 are connected by an unbroken chain of appropriate transitions. Without time explicitly, we can think of $I(s,t)$ as a relation that implies s and t are "spatio-temporally or causally continuous in the appropriate way." We expect I to be reflexive and symmetric, probably transitive (if continuity from A to B and B to C, then A to C), though transitivity might fail if "integrity" is only defined for short intervals and needs something else for long ones. But likely we enforce it as an equivalence relation too (maybe just one big equivalence class per tag if no branching). Actually, if we incorporate branching, I could still be an equivalence (two points in the object's history are connected through it), but branch splits complicate it – that's why we separate NB. In essence, $I(s,t)$ says s and t belong to one continuous life/trajectory. In models, one might define $I(s,t)$ in terms of the existence of some path from s to t along temporal stages that satisfy certain coherence (like how in an Aristotle/Aquinas sense, an object's matter or form flows). For formal logic, we'll treat I axiomatically, requiring it to hold in obvious cases and disallow breaks as needed. We will also allow I to depend on sortal – e.g. our default policy might be metric continuity: the positions of s and t in spacetime are close enough, etc., but we won't formalize that fully here. We may however impose that if $I(s,t)$ fails, identity fails, etc. And if the scenario requires, one can parametrize I by type – e.g. $I_{human}(s,t)$ meaning continuity as a human. - **No-Branching policy:** A predicate $NB_π(s, t)$, sort $S \times S$ relative to a tag (we may treat it as part of the R relation or separate). $NB_π(s,t)$ means *"under tag π, between s and t, no branching has occurred."* More concretely, if s and t are two stages or instances on the same worldline (same π), NB holds if the segment of history from s to t did not split into two distinct existents at any point. If there was a split, NB would fail. We might formalize $NB_π(s,t)$ as: for all u, if $R_π(s,u)$ and $R_π(t,u)$ and $I(s,u)$ and $I(t,u)$, then (u is either in the past of both or after both or something) – basically ensure uniqueness. A simpler way: $NB_π(s,t)$ is true iff there is exactly one maximal chain of I-connected $π$-tagged objects that includes both s and t. If a branch occurred, then there would be two distinct such chains after the branch, thus s and t would not both be on a single chain (except trivial case at branch point perhaps). We will treat $NB_π$ as a needed condition in identity: meaning if the origin splits, identity cannot be asserted beyond that point. For most secular scenarios, one can assume $NB_π$ is always true unless explicitly given a fission. In theological terms, since divine persons have different π, NB is moot there (no one else shares Father's π anyway, so trivial). - **Identity:** We include the usual equality symbol "=" for each sort (or at least for S sort, but we could have for N and Π too in the meta language). Classical logic normally has one = but in many-sorted one can allow different = per sort. We will have equality on S (written $s = t$ meaning they are the same supposit). However, we will *not* treat this as primitive in our theory – we will not assume the standard Leibniz axioms for it globally. Instead, we will define a new binary predicate `Id(s,t)` (or use $=$ in the surface language carefully) meaning *"licensed identity of s and t holds."* In formalism, we might not need a new symbol if we just define $s=t$ as an abbreviation. But to avoid confusion, think of `Id(x,y)` as a formula that can be true or false and is *not* just the trivial built-in identity. We will in fact define `Id` in terms of $R, Con, I, NB`. Ultimately we want something like:`Id(s,t) := ∃π [R_π(s,t) ∧ Con(s)=Con(t) ∧ I(s,t) ∧ NB_π(s,t)]`. This will be the formal definition of identity from causes. Once we have that, we might restrict the usage of the built-in "=" to either not appear at all for $S$ (treat it as just`Id`), or to appear only in extensional positions as a shorthand. Typically, one can retain the built-in $=$ for S but then the logic's rules must be modified so that $=$ is not freely substitutable – effectively demoting it. It may be easier to not use built-in $=$ for supposit sort at all, and only use`Id(s,t)`as the identity test. In the meta-logic we'll still say things like "for any object x, $x = x$," but in the object logic of S we'll use`Id`.

Given these symbols, the **well-formed formulas** are built as usual with logical connectives ($\neg, \land, \lor, \to, \leftrightarrow$) and quantifiers ($\forall, \exists$) sorted appropriately. We'll have to be careful that, for example, if we quantify over S, we only plug those variables into $P(-,-)$ or $I(-,-)$ etc. The sorts enforce that.

We will also allow some **macro notations**: - We write $s ≡_π t$ for $R_π(s,t)$ (just a more readable infix for provenance equivalence). - We might write $s \equiv t$ if we mean $\exists \pi\, R_π(s,t)$, i.e. they have some common provenance (without specifying which). But usually we do specify π if important. - We write $s = t$ in the surface language when we mean $Id(s,t)$ and it's a context where we know that is extensional. In formal proofs, we might avoid using "=" altogether except as a meta-language convenience. - We write Con(s) = Con(t) as shorthand for: $\forall n [K(s,n) \leftrightarrow K(t,n)]$. That asserts the set of natures of s and t coincide. We might also introduce an operator $\mathrm{Con}: S \to \mathcal{P}(N)$ in semantics. But as formulas we can express equality of those sets by mutual inclusion.

To avoid heavy formalism in the main text, we will often speak in the mixed natural-language-plus-symbols style, but the above indicates the rigorous underpinnings. Appendix A will provide the formal axioms, but here we will describe them in prose.

3.2 Identity from causes: formal definition

Now we present the definition of the **licensed identity relation** `Id` within LFI-π. In plain prose:

> *s is identical to t if and only if there exists some origin tag π such that s and t share the same π (i.e. $R_π(s,t)$ holds), s and t have the exact same set of constitutive natures (so for every nature N, $K(s,N) \iff K(t,N)$), s and t satisfy the integrity conditions (they are one continuous being, $I(s,t)$), and there was no branching in the π-line between s and t ($NB_π(s,t)$ holds).*

In symbolic form, as anticipated:

$$ \textsf{Id}(s,t) \;\iff\; \exists \pi \Big(R_{\pi}(s,t)\ \wedge\ [\forall N\, (K(s,N) \leftrightarrow K(t,N))]\ \wedge\ I(s,t)\ \wedge\ NB_{\pi}(s,t)\Big). $$

This single formula captures the core idea: identity is **not primitive** but *defined in terms of more fundamental relations*. It's worth noting that each conjunct plays a critical role: - $R_{\pi}(s,t)$ ensures **common origin** (no identity between things that came from different beginnings). - The K equivalence ensures **same essence/form** (no identity between things that differ in what they are fundamentally). - $I(s,t)$ ensures **continuous existence** (no identity if there's a break or gap in the being). - $NB_{\pi}(s,t)$ ensures **no divergent branch** (no identity if one origin led to two).

One can see this as an enforcement of Leibniz's Law at a deeper level: any violation of these is a discernible difference: - If different origin, they differ by origin property ("x came from cause A, y from cause B"). - If different nature sets, obviously a property difference ("x is of nature N, y is not" for some N). - If not continuous, one might say they differ in a property like "exists continuously from time0 to time1" vs not. - If branching happened, they differ by "has a unique history" property.

Thus, `Id` defined this way will automatically behave such that if `Id(s,t)` then indeed s and t share all such properties (and presumably all properties *period*, but we restrict substitution to extensional ones anyway). We will prove separately that $\textsf{Id}$ so-defined is an equivalence relation on the set of all *stages* of a given individual. Essentially, within any single continuous history (tag), `Id` collapses all the stages into one equivalence class representing the whole individual (the "worldline"). Across different tags, obviously no identification occurs. And at a branch point, the equivalence breaks into picking one branch. This ensures that **on the extensional fragment** (where multiple natures and intensional contexts are not present), `Id` coincides with what you'd normally call equality. In fact, if in a given scenario every object has exactly one nature and no fission occurs, $Id(s,t)$ will reduce to $∃\pi [R_π(s,t) ∧ I(s,t)]$. If we further assume continuity I always holds when $R_π$ holds (i.e. no weird disintegration), then $Id(s,t)$ reduces to $∃\pi\, R_π(s,t)$. If each object gets a unique origin, that further reduces to $τ(s) = τ(t)$ conceptually [19]. If tags are unique to each individual, then $τ(s) = τ(t)$ iff $s=t$ in classical sense (one tag per individual). So in such a simple model, `Id(s,t)` is basically s and t having same unique ID – which is just equality. That's the case in ordinary first-order logic where each element can be "tagged" by itself.

To put it differently, LFI-π is **conservative**: any structure that satisfies these new axioms can be collapsed to a classical structure by forgetting the extra structure (tags, etc.), and you won't derive any new equation or lose any equation in purely first-order terms about = that you couldn't in classical logic. This was proven in earlier work [24] [23]. Intuitively, in an extensional context the extra conditions for identity are automatically met (trivially or vacuously), so `Id(s,t)` holds exactly when classical identity would have, and substitution is allowed there. We'll revisit conservativity in 3.4 with a promise of Appendix proof.

It's important to clarify how `Id` interacts with the sorts and with predication. We will **not** allow mixing sorts in identity statements: we do not even define `Id` for an S and an N (that's ill-typed). So one cannot state something like $Id(s, n)$ meaning "supposit = nature". That sort of statement is simply not in the grammar (and correspondingly, not in theology either: one doesn't say "Christ = human nature", one says "Christ has human nature"). This is how we *forbid category errors by syntax*. Similarly, we won't identify a tag with a supposit or such. The only identifications to consider are S-S, N-N, $Π$-$Π$ perhaps. For N-N, one could ask: do we allow identity between natures? Possibly yes, two nature constants could be identical if they mean the same nature (like *Homo sapiens* vs *human* might be synonyms). But typically, we treat nature symbols as distinct if they are distinct concepts. We could, if needed, add axioms that some N_1 equals N_2 meaning they are the same kind of nature. But that's usually not necessary unless we have synonyms. We can likely leave that to the model (if a model wants to interpret two constants as the same, fine). Since the question of identity of natures is rarely needed to discuss (except saying something like "Divine nature is unique" which one can say $\neg ∃n≠\textit{Divine}: P(Person, n) ∧ n$ has some property – anyway). Similarly, identity of tags we might treat as basic equality on tags (two tags can be equal if they represent the same origin – though in models usually tags are just abstract IDs so equal only to themselves). We won't focus on those; the key is identity on supposits, which we have defined.

3.3 Rules in prose: introduction, substitution, aspect discipline, no global Leibniz

We next outline the **inference rules** and logical principles that govern this system. In natural deduction or sequent calculus style, we have to modify or add a few rules compared to ordinary first-order logic with identity. We'll describe them:

- **Id-Introduction (Id-Intro):** This rule allows one to conclude $Id(s,t)$ (or $s=t$ in the surface sense) when certain conditions are met. Of course, one way is if s and t are syntactically the same variable (reflexivity), we trivially have $Id(s,s)$ since all conditions hold trivially: same tag (itself), same Con (identical obviously), integrity (reflexive), no branching (reflexive case trivial). So reflexivity $Id(s,s)$ is always provable, just like $s=s$ is in classical logic [20]. But the interesting introduction is if we know or can establish each of the required conditions. For example, in a proof, if we have: (a) a specific tag π such that $R_π(s,t)$ holds, (b) for every nature N we have $K(s,N) \iff K(t,N)$ (in practice, often one would show each $K(s,N)$ implies $K(t,N)$ and vice versa if a finite list, or appeal to some given equivalence of their constitutions), (c) $I(s,t)$ is known (perhaps given or deduced from continuity assumptions), and (d) $NB_π(s,t)$ holds (likely one would argue this from no interference or an assumption that π had unique continuation between those points) – then we can conclude $Id(s,t)$. This is akin to constructing an identity by showing all necessary evidence. In everyday terms, if you want to prove two occurrences are the same person, you'd produce evidence they were born from the same mother at the same time, have the same fingerprints (analogy to Con set, same attributes), and that the person went continuously from one place to the other (integrity), and that there was no twin or split (no branching). Under those circumstances, you'd be justified in saying "it's one person". That's Id-introduction in narrative form.

- **Licensed Substitution (Id-Elimination or Substitution):** In classical logic, from $x=y$ one can substitute y for x in any formula. In LFI-π, we **restrict substitution to declared extensional positions**. What is an extensional position? These are contexts where the meaning of a formula doesn't change if you swap a term with an identical term. Typically, ordinary atomic formulas like $P(s,N)$, $K(s,N)$, $I(s,t)$, etc., are extensional in the s,t positions (they're just relations on those terms). More generally, any formula built without intensional operators (like modal or belief contexts) is extensional in all term positions. So inside such a formula, if we know $Id(a,b)$, we can replace a with b. However, if a context is marked as intensional (say we had an operator $Believes(Lois, \_)$ or a modal $\square$), then we restrict that one cannot substitute inside it unless additional conditions are met (like the provenance in the intensional context as well). Since our core calculus doesn't explicitly include modal or belief operators, this mostly applies when using the logic in an enriched form. But we maintain the principle: **only substitute where it's safe**. This is enforced in practice by not having a blanket Leibniz's Law axiom, but rather proving case-by-case that certain contexts allow it. For example, we can prove as a theorem schema: if $Id(x,y)$ and $\phi(x)$ is a formula with x only in extensional positions, then $\phi(y)$ can be concluded. This we call **Licensed Substitution**. Because our logic internalizes aspect and type constraints, many potential problematic substitutions are simply ill-typed (e.g. substituting a whole term for a part of a term? Not possible here; or substituting across nature vs person? Cannot). So the remaining risk is intensional contexts which we are not fully formalizing here. So one can simplify: in LFI-π as presented (without extra intensional ops), if $Id(s,t)$ is derived, one can treat s and t as interchangeable in all atomic predicates and thus all usual formulas. This gives us reflexivity, symmetry, transitivity of Id easily:

- Reflexivity: $Id(x,x)$ (by intro or axiom).
- Symmetry: from $Id(x,y)$ we get $Id(y,x)$ because all the conditions are symmetric (we can prove that in the calculus).

- Transitivity: if $Id(x,y)$ and $Id(y,z)$, then because y can be substituted for x in $Id(y,z)$ we get $Id(x,z)$, or more directly one can show the existence of a common tag and continuity from x to z. We will have an axiom or derived rule for transitivity. Essentially, sharing the same tag and each with continuity to y implies x,z share that tag and continuity via y, so $I(x,z)$ and no branch (if a branch happened in one half, it would break Id in that half), plus if Con sets matched x,y and y,z, by substitution they match x,z. So yes, $Id(x,z)$. All that can be done formally [21] [24]. So `Id` is an equivalence relation on each sort domain where it applies.

- **Aspect Discipline:** This is not a single rule but a general restriction: the logic distinguishes "aspects" like supposit vs nature in statements to prevent illegitimate inference. One manifestation is the typed equality we already have (no mixing sorts). Another is an explicit rule: you cannot conclude or assume an identity or predication across aspect domains. For example, we do not allow any rule that would conclude $K(s,n)$ from $P(s,n)$ or vice versa without justification. In fact, we will include an axiom that $K(s,n) \to P(s,n)$ (if something is constituted by a nature, it truly is of that nature) 【4.3】, but not the converse (a supposit might truly be of a nature but perhaps we don't list it as constitutive? Actually if it's truly of that nature, probably it should be constitutive; maybe an exotic case: Christ is truly mortal (predicate), but mortality is not a separate nature in his Con set, it follows from human nature; anyway, our system simplifies by saying predication can hold either because of constitution or because of some accidental property. In any case, we keep track).

The main aspect discipline arises in theological usage: if we have $P(s, Divine)$ and $P(s, Human)$, we never derive a contradiction because any property we apply, we have to specify which nature's aspect it belongs to. For instance, if F is a predicate meant for divine nature (like "eternal"), we can have a rule that if $P(s, Divine)$ then $F(s)$ means actually F of that nature is true. But if we have a predicate G for human nature (like "suffers"), from $P(s, Human)$ we get $G(s)$. Now a naive logic might then allow "$F(s)$ and $G(s)$, and since s is one individual, maybe derive something silly like F and G apply to the same subject in a contradictory way. But in our disciplined approach, we consider that F and G might be incompatible if thought of as about the same nature, but here they are about different aspects of s. We enforce that by implicit context or sorted predicates. Possibly we will have separate predicates for divine operations vs human operations, etc., or a sort system where one cannot apply a purely divine property to a human nature. However, since one supposit has both natures, we have to allow s in both contexts. The discipline is enforced more semantically by how we interpret the statements, and by adding "guards" in formal derivations (like a rule: from $P(s,N)$ and a property F that applies to N, infer $F(s,N)$ where we attach the nature as parameter, to keep track). Because the user requested minimal symbolism in body, we won't formalize that fully here, but the promise (from Section 4.3 and 7.2) is that our calculus prevents mixing up predications without awareness of which nature they belong to【4.3】. Concretely: it will never prove something like "As God, Christ died" without a caveat that "died" is only true according to human nature. It will allow "Christ (who is God) died" in an ordinary language sense because that can be paraphrased "The Person who is God died in his humanity," which our formalism can express safely (with `Id` linking the person and nature roles properly [41]). In summary, the aspect discipline means **no cross-sort identification, and careful tracking of which sort each predicate pertains to**. - **No global Leibniz's Law axiom:** In classical FOL, one often has an axiom scheme: $x=y \to (Φ(x)\to Φ(y))$ for any formula $Φ$ or a specific restricted version. We do **not** include a blanket scheme like that for Id, because of intensional contexts. Instead, we rely on the derived rules of substitution in safe contexts. This omission is crucial: it is what blocks things like substituting "the clay" for "the statue" in a modal context like "necessarily, the statue is statue-shaped" (the clay is not necessarily statue-shaped since it could have been mushed, so that would break). By not having the law, the move is not allowed unless we justify it with additional premises like NB and continuity under modality. In essence, we built the restrictions into the definition of Id (so if a property fails to be shared, Id wouldn't be true in the first place) and into usage rules (so even if Id is true, we don't blindly push it into every context).

One might ask: is equality now a logical notion or a theory-specific notion? It's a theory-specific defined predicate. That means some general meta-theorems (like existence of certain normal forms) might need rechecking, but since we ensure Id acts like = in classical fragment, completeness and soundness relative to classical logic hold in that fragment [64]. It's basically a conservative conservative extension with more expressive power but same consistency strength.

Finally, we mention that in the background logic, we still have classical logic for non-identity stuff. E.g. we still use all quantifier rules, connectives, etc. The only deviance was the identity handling and the additional axioms around R, I, etc. We also likely add axioms to enforce how these new predicates behave: - $R_π(s,s)$, $R_π(s,t)\to R_π(t,s)$, and if $R_π(s,t)$ and $R_π(t,u)$ then $R_π(s,u)$ (so each π defines an equivalence on S) [21]. - Uniqueness of tag: maybe an axiom "If $R_π(s,t)$ and $R_{\rho}(s,t)$ then $\pi=\rho$". In other words, an ordered pair of supposits can't share two different tags unless those tags are equal. This enforces that a given pair of things either have exactly one common origin or if multiple, those must effectively be the same origin. This prevents weird overlapping multi-origin scenarios (except maybe Christ who has divine origin and human origin if we considered two aspects; but in our approach, Christ's π might be just his human origin, his divine nature isn't via a creaturely cause, so we treat divine origin separate; anyway skip detail). - I equivalence: reflexive, symmetric, transitive on S (or at least on those connected by R). - If $R_π(s,t)$ and $I(s,t)$ and $R_π(s,u)$ and $I(s,u)$ and $R_π(t,u)$ fails, that would indicate branching. We probably formalize $NB_π(s,t)$ as: for every u,v such that $R_π(s,u), R_π(s,v), I(s,u), I(s,v)$ and we have $I(u,v)$, then either $I(u,t)$ and $I(v,t)$ or u or v is outside some range. Actually, might be easier: define a notion of time ordering on indices and require that for all times between s and t, only one object with tag π exists. But formalizing that fully requires introducing time explicitly. Another approach is to treat the domain as stages with a partial order and enforce each tag's set of stages is linear. We might just state: $NB_π(s,t)$ fails if and only if there exist $u\neq v$ such that both $Id(s,u)$ and $Id(s,v)$ would hold if not for branching (i.e. they share π and continuity up to s but then diverged). Perhaps we can leave NB somewhat intuitive here, as its main use is in the case studies to say "two futures -> not identical". - We might have an axiom that says something like: if $Id(s,t)$ and $P(s,N)$ then $P(t,N)$ (that's a restricted Leibniz for P). Similarly for K, I and even for R (if identity holds they share all same tags which they do by construction). We will ensure that any predicate representing an extensional property is preserved by Id. Many of those follow from the def: e.g. if $Id(s,t)$, they share tag by existence of some π, and likely the way we define it ensures that's the only tag either has – maybe we assume one principal tag per supposit, so we can deduce if they share one tag and tags are unique, they share all (the one). For nature, Con set equality is built in. For an arbitrary extensional predicate F, we might make it explicitly or just assume any primitive predicate we add will be tag- or nature-dependent so that Id entails it. We could restrict the language that any new predicate must either be a function of s's nature or something that does not violate identity if Id holds, to maintain coherence. Those are design choices – akin to saying the only properties that matter are those that depend on origin, nature, continuity, or combos thereof. If someone introduced a property that doesn't depend on those but could differ for identicals, that would break the idea of Id capturing all discernibles. But by design, we consider those conditions to cover what normally are thought to individuate.

Summing up rules: **Id-intro**: prove all conditions, infer Id. **Id-elim (substitution)**: if Id and context extensional, substitute. **Reflexivity**: always $Id(x,x)$. No unconstrained Leibniz's Law. **Constitution rules**: if $K(s,N)$ then $P(s,N)$, and likely if $K(s,N)$ and $K(s,M)$ and maybe some exclusivity if needed (maybe one supposit can have multiple though). Possibly an axiom "for each supposit s, and each nature N, either K(s,N) or ¬K(s,N)" if we want classical logic on K, but maybe not – better to leave it open because one might not have a closed list of natures for a supposit. **Typing rules**: you can only put S terms in S slots, etc. That's enforced by grammar.

3.4 Conservativity and classical shadow

We have emphasized that LFI-π is **conservative** over FOL with =. Let's articulate why and in what sense: - Any formula in the language of ordinary first-order logic (without the new symbols P, K, R, I, etc., just using =) that was a theorem there, will still be valid (provable) in our system *when restricted to extensional scenarios*. More formally, if you have a formula $Φ(x_1,...,x_n)$ in pure logic with = and you interpret = as Id, then if it was logically valid classically, it's derivable in LFI-π *provided that any occurrence of Id in it is only applied to terms that indeed satisfy the licensing conditions in the model*. Perhaps a simpler statement: "On any model of LFI-π, the equality relation (as defined by Id) on extensional objects is an equivalence and respects all first-order indiscernibility properties." Hence no contradiction or new equation can be derived just by the presence of the extra machinery. We proved in prior work that no "pure =" formula that wasn't tautologically true becomes provable, nor one that was consistent becomes inconsistent [24]. - Conversely, any identity statement that holds in the classical sense (meaning just that the two referents are the same object in the intended model) will be provable to hold via Id, because if they're the same, trivially they share origin, nature, etc. So any true identity can be shown with our stricter criteria (the real object meets them because it's literally one object). - The main difference is: some things that classically would be *either not expressible or false* become *expressible and possibly true* in our system. For example, classically one might just say "Lumpl = Goliath or not"; in our system we can say more: "Lumpl ≠ Goliath, but Lumpl ◇ Clay and Goliath ◇ Clay and they share matter continuity etc." We haven't lost classical equality (we still can talk about distinctness of lumpl and goliath and all), we just added a framework that elaborates that scenario more richly. - In any purely extensional context, we can introduce a macro: define $s *={ext} t$ to mean $Id(s,t)$ when s and t have no special intensional properties or multi-nature. Then we can show that $*= t$ then for any predicate F (that is extensional) $F(s)\iff F(t)$. So the usual Indiscernibility of Identicals holds in that fragment [21] [24]. That essentially shows classical logic is a }$ behaves exactly like classical =: reflexive, sym, trans, and if $s =_{ext*sub-theory* of ours (any classical proof about extensional stuff is mimicked). - The conservative extension lemma would be: if a sentence ϕ in the language without new symbols is provable in LFI-π, then ϕ was already logically valid in classical logic with identity. And if ϕ is consistent classically, then adding our axioms doesn't make it inconsistent. We have not explicitly given the formal proof here, but we refer to earlier completeness arguments [65] [64]. The key to proving it is usually constructing a structure: given any classical model, we can expand it to a model of LFI-π by assigning each element a tag (like itself or a surrogate ID) and letting I be total (everyone continuous with themselves trivially), NB trivial since no splitting in a static model, etc., and nature sort maybe each element has a nature "Self". That expanded structure satisfies our axioms and in it Id coincides with original identity. Thus any formula without new relations that was true in the classical model remains true in the expanded. Therefore no new entailments appear. - We deliberately ensure that our additional rules (like substitution) do not allow deriving false conclusions in classical terms. For instance, we wouldn't accidentally prove something like $\exists x ¬(x=x)$ (which is unsatisfiable) because reflexivity of Id prevents that and we've not broken consistency.

In summary, **where our identity applies, it acts like classical identity**, and **where classical identity would lead to paradox, our identity simply doesn't apply** (stays silent, because conditions fail). For a simple secular example: classical logic cannot say "these two are partly the same and partly not" – either they're identical or not. Our logic can say "these two objects are not identical (because Con differs), yet they share matter and origin," which classical logic couldn't express except in second-order or vague terms. So we refine the logical expressiveness.

We might dub classical equality as a **shadow** inside our logic: If we define a context to be "extensional" if a certain set of conditions hold (like same tag and no intensional ops), then within that context we can use the symbol "=" to mean Id and all usual reasoning goes through. In fact, we do treat "=" as just a notation for Id in extensional places【4.1】. Readers of our paper will see statements like "Let $x=y$ for extensional x,y," and they can interpret it just as normal equality because indeed it will behave so.

To connect to known ideas, it's a bit reminiscent of how in **Homotopy Type Theory (HoTT)** they have an identity type that can be non-trivial (there can be multiple identifications or higher paths). Our identity is not as radical (we still treat identity as a proposition that either holds or not, not many paths), but we share a motivation of *replacing a too-strong notion of equality with a more structured one.* We maintain equivalence and congruence in normal cases, but allow distinctions where needed. The novelty is we did it in first-order logical terms with extra predicates rather than type theoretic or second-order apparatus.

One last note: Because we eliminated primitive equality, one might worry "how do we talk about how many objects there are, or functions, etc.?" In many-sorted logic we often allow equality in each sort. We have something similar: we do allow equality on the tag sort and nature sort. Those are fine as primitive if needed, or we can define them similarly. But equality on supposit sort is defined. If we have function symbols on S, we have to clarify how they interact: presumably if $Id(x,y)$ then $f(x)$ and $f(y)$ should also be identified. We'd ensure functions produce same tag and nature if inputs do, so results come out identical if inputs identical – making Id a congruence relation on S domain operations [21] [24]. We can add axioms: $Id(x,y) \to Id(f(x), f(y))$. Then one can mod out by Id and get a well-defined function on equivalence classes. This is standard for any equivalence extension.

Thus, the formal core of LFI-π is laid out: a richer language, a definition of identity via cause and profile, rules that restrict substitution to safe contexts, and a conservation of classical reasoning in its appropriate realm. With the core theory in hand, we proceed next to Section 4, where we introduce a more **readable surface layer (PEI-L)** to use this logic without drowning in symbols, showing how an ordinary user can write and prove statements clearly using our notations and where we hide the licensing details under the hood except when needed.

4. Readability layer: PEI-L surface

4.1 Notation from PEI: typed predication, constitution, provenance

The formalism introduced in Section 3, while precise, can be cumbersome to use directly. To make the system **reader-friendly**, we employ a surface language inspired by the original PEI (*Provenance–Esse–Integrity*) framework [26]. This surface syntax, which we call **PEI-L**, uses intuitive mathematical symbols for the key relations, allowing us to write statements in a form that closely mirrors natural philosophical language. The guiding idea is to present formulas in a way that a philosopher might write on paper, while knowing that in the background these are shorthand for the more complex formal relations of LFI-π.

Here are the main notational conventions of PEI-L: - **Predication:** Instead of writing $P(s, N)$ for "s has nature N", we write $$s \in N.$$ This suggests set-membership intuitively (the individual belongs to the class of things of nature N) and was used in prior PEI papers [26]. For example, `Socrates ∈ Human` means Socrates is human (i.e. Socrates has the human nature). In a theological context, `Christ ∈ Divine` and `Christ ∈ Human` would be used to say Christ is divine and human, respectively. The symbol "∈" is just suggestive; it does not mean we treat natures as sets in a literal sense, but it aligns with the idea of an individual falling under a nature or kind. - **Constitution:** We write $$s \◇ N$$ (some texts may use $\diamond$, but we'll use a specific diamond or star symbol) to denote $K(s,N)$, meaning nature N is part of the constitutive essence of s. This notation was likewise used in PEI [66]. For example, `Christ ◇ Human` and `Christ ◇ Divine` indicate Christ is constituted under both humanity and divinity. In more everyday terms, if we had an object that is, say, a statue made of clay, we might say `Statue ◇ Clay_nature` and `Statue ∈ Statue_nature`. The difference between "◇" and "∈" is subtle but often in practice a being's nature in predication will coincide with a constitutive nature unless we're considering accidental properties. We use "◇" especially when an individual has multiple natures to keep track of them in a set (Con(s)), or when we want to emphasize aspect ("s qua having nature N"). - **Identity (licensed):** In the body text, when discussing extensional situations, we will sometimes use the ordinary equals sign "=" between terms, with the understanding that **within this text, "=" is not a primitive logical identity but a shorthand for our defined identity in the appropriate context**【4.1】. We will ensure whenever we write $x = y$ that it's either (a) in an extensional context where our rules allow it to behave classically, or (b) we immediately clarify it as meaning $Id(x,y)$ in LFI-π. Essentially, the symbol "=" in our exposition is *macro sugar* for "all the licensing conditions hold, thus we assert identity". For instance, we might write "If $a = b$ (meaning they share all requisite causes and continuity)…" or in proofs "since $R_π(a,b)$, $\mathrm{Con}(a)=\mathrm{Con}(b)$, etc., we have $a = b$." By doing this, we don't burden the reader with $Id(a,b)$ and existential quantifiers every time – we state it in plain math equality. However, behind the scenes the meaning is always the defined one, not an assumed absolute identity. We caution that whenever "=" appears in a formula in this paper's formal sections, it really means the defined identity *in an allowed position*. We avoid using "=" in tricky intensional examples to prevent confusion. We also do not use "=" between an S and an N (sorted distinction forbids that) – so one will not see something like "Christ = Divine" in our formal prose; instead we say "Christ ∈ Divine" or "Christ is (a) God" in words. - **Provenance relation:** We often simply write something like "x and y have the same origin π" in prose. If needed formally, we write $x ≡_π y$ to denote $R_π(x,y)$ [28]. If the specific tag is unimportant or understood, we might just say "share origin" or "co-provenance". In the surface proofs, we might use phrases like "by origin" or "due to same π" to indicate using the fact of common provenance. For example: "Since the Ship on day 0 and the Ship on day 100 share the same tag (origin), we label them with π and proceed…" We may not use a special symbol for existence of some tag – often it's clear from context ("there exists some π such that …"). - **Integrity and continuity:** We will use plain language to say "continuous from s to t" or "no break between s and t" to invoke $I(s,t)$. If needed, we might write something like $s \rightsquigarrow t$ as a visual of a continuous path (this symbol not standard, but just for exposition). However, usually this is in proofs where we say "because there is a continuous path (integrity) from s to t, we have …". So we rely on words rather than a symbol for I. - **No branching:** We again use English: "assuming no branching occurred" or "without branching" or "under the no-branch condition NB." If we did use a symbol, maybe something like "$\nwarrow$" (no fork) might be intuitive, but we won't introduce a new symbol unnecessarily. It's typically a condition mentioned as a side note: e.g. "object A remained one (no branching) through the process, therefore…".

To illustrate, consider how we might rewrite the core identity definition in this surface notation:

"**$x = y$ iff $\exists \pi$ such that $x ≡_π y$, $x \◇ N \Leftrightarrow y \◇ N$ for every nature N, x and y are continuously connected (no interruption in being), and there is no branching of π between them.**"

This reads much more naturally than the raw formula, yet it conveys the same content【4.1】. In a running argument, one might break it up: "Suppose x and y come from the same source π and neither gains or loses any nature the other doesn't have, and suppose furthermore that x develops into y with no break or branching. Then we conclude $x=y$."

The PEI surface notation also carries **suggestive semantics**: - $s ∈ N$ suggests "s is an instance of N" (like set membership). - $s ◇ N$ suggests "s has N as a component of its being" (one might say "s is an N-composed entity"). - $≡_π$ suggests an equivalence by tag. - We will continue to use standard logical connectives in the usual way. - We keep the overall page free of too many quantifiers by often stating things in words ("for all natures N such that…, …").

One should note that **the symbol "=" will never be used across sorts**. We do not write "Human = Animal" or "Christ's human nature = His divine nature" – those are false in meaning and invalid to even formulate under sorted usage. Instead, equality is only between two supposits or two natures of the same sort. If we want to say two natures are the same kind (e.g. two different names for the same nature concept), we might just state that in words, since it rarely comes up except maybe "rational animal nature = human nature” in philosophy, which we could express as an identity of nature constants. But that’s more of a theoretical identity (essence definitions) rather than individuals.

Importantly, **we do not search for or introduce any new symbols beyond those** unless needed. The user specifically said not to go hunting for images, and similarly we aren’t hunting for new fancy notation beyond what’s given by PEI tradition【Guideline】. We have them: ∈, ◇, ≡, =. That suffices.

4.2 Quick sequent patterns for ease of use

For readers and users of this system (especially those proving theorems within it), we provide some **template patterns** or inference moves that commonly occur, so they can be recognized and not re-derived each time. These patterns essentially show how to use the PEI-L notation to streamline proofs:

- **Extensional substitution pattern:** If you want to substitute a for b in some formula, you show $a = b$ in an extensional context. In practice, this means:
- Show $a ≡_π b$ for some specific π (i.e. identify a common origin).
- List the natures of a and show b has exactly those (or vice versa). In text, one might say “$Con(a) = Con(b)$ because...”.
- Assert continuity: “from a to b there’s no loss of identity (they form one uninterrupted entity)”.
- State no branching: often by context, e.g. “assuming uniqueness of this line of descent”.

- Conclude $a = b$. After that, one is free to replace a with b anywhere (in formulas about nature membership, properties, etc.). So a sequent might look like:

  From $a ≡_π b$, and $\forall N (a \in N \iff b \in N)$, and continuity($a,b$), and no-branching($π$ on a,b), infer $a = b$.

  Then from $P(a)$ deduce $P(b)$ for any predicate P that doesn’t involve intensional context.

This is basically how one *applies* Id-elimination in practice.

- **Proof of identity by cases pattern:** Sometimes one must prove two things are identical by considering various scenarios (like either the object moved or a new one replaced it, etc.). In such a case, one can break into cases and in each case, verify the conditions. For example, to prove that after a plank replacement the ship is still “the same” by our criteria, one might do induction on planks: (i) Base: 0 planks replaced, trivial identity. (ii) Step: assume true for n replacements; for the $n+1$th, check that the tag remains the same (the ship’s continuity doesn’t break because you replaced one plank), natures remain (it’s still a ship made of wood, presumably Con same), and no branching (the old plank is discarded not making a second ship). Therefore by induction all conditions hold after all replacements, concluding identity persists. This pattern shows *how to integrate our identity conditions into proofs by induction or case analysis.* Each step you ensure provenancial continuity and no new natures. That is straightforward to articulate in our notation: “The Ship with plank 5 replaced ≡ same tag as original (since we keep same hull, presumably), Con unchanged (it’s still wood hull, structure intact), integrity maintained (we do it in dock gradually, not destroying it), no branching (no duplicate ship). So by induction, the final ship after all replacements = the original ship.” We implicitly used a sequent like above at each step.

- **Guarded mixed reasoning pattern:** If one wants to combine statements about different aspects (like using a property of Christ's human nature and a property of His divine nature together), one must **guard** the reasoning with explicit mention of which aspect is invoked. For instance: "Christ ∈ Human implies Christ has a passible (capable of suffering) nature; Christ ∈ Divine implies Christ has an impassible nature. However, because *passibility* or *impassibility* are predicated of different natures (the first of the human, the second of the divine), there is no contradiction in saying Christ (one person) is both passible (in his humanity) and impassible (in his divinity)." The pattern here is:

  To show X and Y are consistent properties of the same subject s, show there are N_1, N_2 such that $s \in N_1$ and $s \in N_2$, and X applies to N_1-aspected predications of s while Y applies to N_2-aspected predications of s. In formal steps: 1. $s \in N_1 \land s \in N_2$ (with $N_1 \neq N_2$). 2. X pertains only to nature N_1 (for example: "mortal" might be defined as "having an animal body that can die" – applicable to human nature). 3. Y pertains only to nature N_2 ("immortal" could mean "having divine life"). 4. Therefore $X(s)$ is true in virtue of $s \in N_1$, and $Y(s)$ is true in virtue of $s \in N_2$. We mark this by perhaps writing $X(s|N_1)$ and $Y(s|N_2)$ to indicate the aspect. 5. No logical clash because the logic will not equate $X(s|N_1)$ with ¬$X(s|N_1)$ or such – they are different predicates fundamentally (like $Mortal_{Human}(s)$ vs $Immortal_{Divine}(s)$).

We might not formalize it exactly like that in text, but that's the reasoning pattern. The result is a *communication of idioms* scenario solved by aspect separation. The logic's discipline ensures one cannot combine those into a single contradictory predicate because they live in different sortal contexts.

- **Conservativity usage pattern:** If working in a domain where none of the fancy stuff is needed (say a plain math proof by induction on natural numbers), the user can basically ignore all the LFI machinery. One just uses $=$ as usual. Under the hood, each number has a unique tag, integrity is trivial (numbers don't change), Con could be (Number) for all, etc., so everything goes through. This pattern is "just do what you'd normally do; the logic won't interfere." So a mathematician can apply all normal equational reasoning and never worry about R or I because those axioms are quietly satisfied without fuss. We have thus patterns where a user might consciously "set intensional knobs to off" and proceed classically. In our paper, we might illustrate that by doing an example of a secular puzzle solution and highlight that in a straightforward scenario it looks exactly like a solution in classical logic plus perhaps a mention of cause tracking. This assures readers the system doesn't break normal reasoning. It's helpful because often new logics restrict something and require new unnatural proof steps – our claim is in the extensional realm, proofs look the same.

By providing such patterns in explanatory sections (maybe in a methodological aside or a footnote), we aim to make it *easy for readers to follow or reconstruct proofs without being bogged down*. Essentially, we hide the overhead. The formal calculus might require four facts to use an identity, but the author in text might say "clearly, given the circumstances, these two are identical (same origin, form, continuous, etc.) so we substitute one for the other." That one sentence compresses a chunk of formal steps, which is fine since the readers can fill in if needed and we cited the reasons in parentheses.

4.3 Safety by design: preventing category and aspect errors

One of the chief advantages of our PEI-L design is that it **prevents certain kinds of logical mistakes by construction**. We call this "safety by design." We highlight a few specific safety features: - **No person-to-**

nature identities: As emphasized, the system simply does not allow one to equate a supposit with a nature. For example, one cannot prove or even state something like “Jesus’s divine nature = Jesus’s human nature” (which would be heretical and nonsensical). The sorts forbid it. If someone tries to treat a nature as an individual or vice versa, it’s a type error. This mirrors doctrinal correctness: a person is not a nature; Christ is not identical to either of His natures – He *has* them. Our sorted logic enforces that metaphysical distinction at the grammatical level. So certain heretical or paradoxical formulations are literally unformulable. This is a feature: it forces a user to phrase things properly (as predications rather than identities). The result is that many potential logical contradictions (like “God is mortal” as an identity of predicates) cannot even get off the ground. - **Predications stay within their aspect:** This means when you assert a predicate of a subject, it is always understood under the appropriate nature that makes that predicate meaningful. In practice, we either treat predicates as already aspect-specific (e.g. “dies” might implicitly apply to something considered in its animal nature), or we ensure to qualify it if needed. The logic, by requiring an $s ∈ N$ for certain N-specific predicates, ensures you can’t say s dies unless $s ∈ MortalNature$ or something. In our theoretical formalization, we might have certain predicates only applicable to certain sorts or require certain Con membership to even mention. For example, define a unary predicate $Dies(x)$ but really it has a precondition “x has a nature capable of death.” We could enforce that by an axiom: $Dies(x) \to (\exists N (x ∈ N \land MortalNature(N)))$. But simpler: have sorted predicates like $Dies_{Human}(h)$ which only accepts human sort. However, since we kept one sort S for persons and just differentiate by Con sets, the enforcement is more semantic: one shouldn’t assert $Dies(God)$ except as shorthand for “the Person who is God died according to human nature.” And our logic would require that human nature presence is there for it to be true. So a sequent: $s ∈ Human ⊢ Dies(s)$ might be allowed if being human is known to entail mortality, whereas without $s ∈ Human$ one cannot derive $Dies(s)$. That’s one reason to separate predication and constitution: it’s easy to attach rules like “if $Mortal \in Con(Nature)$ and $s ∈ Nature$ then s can die.” Without that typed structure, classical logic might let you attempt to apply “dies” to anything and you have to add an ad hoc premise “only humans can die.” We integrate that naturally by sorts or axioms. - **Constitution strengthens predication:** We included or will include an axiom that if $s ◇ N$ then $s ∈ N$【4.3】. This ensures no weird case where something “has” a nature in its constitution but we don’t consider it truly of that nature. In other words, if an entity’s Con set contains N, it genuinely is an instance of N. This is safety because it prevents misuse of $◇$ as some kind of “potential nature” that isn’t actual. We want to avoid someone saying e.g. “Christ ◇ Divine but Christ ∉ Divine” (that would be contradictory to our design – if He’s constituted by the divine nature, He is divine). So we enforce consistency between $◇$ and $∈$. Conversely, not every predication needs a $◇$ (accidental or lesser categories might be predicated without being in Con). But any $◇$ is a real predication. So one won’t erroneously treat something as fundamental nature without it being so. - **Provenance is independent evidence:** By introducing origin tags explicitly, we avoid relying solely on maybe ambiguous qualitative evidence for identity. In practice, this means the logic won’t let you conclude identity just from “they have all the same properties” unless one of those properties or separate premise indicates common origin. This is a safety measure philosophically: it aligns with the view that just because two things are qualitatively identical doesn’t mean they are the same (could be identical twins, etc.). Classical logic, with the identity of indiscernibles, often leaps that if truly no property differs it must be same – which can be problematic if one thinks maybe haecceity or origin isn’t captured in the listed properties. Our system explicitly includes origin as a condition, highlighting it as separate evidence. So logically, one cannot equate two entities unless something like $R_π(x,y)$ is given. This prevents mistakes like concluding “theseus’s ship rebuilt from planks = theseus’s sailing ship” just by seeing they share all planks at some different times – our logic would catch that they do *not* share the same continuous history (we have to attach tags differently). So requiring a $R_π$ condition ensures identity proofs always mention the historical connection, not just snapshot similarity. That’s a methodological safety: it forces acknowledging how the sameness is established (through cause, not just through accidental similarity). This is something philosophers have emphasized (e.g. "same matter, same form, but are they same *entity*? need cause link" – we have that). - **No accidental collapse of distinctions:** By "accidental collapse" I mean the logic will not inadvertently identify two things just because the user forgot a branch or an aspect. For example, if an author carelessly tried to prove Father = Son by listing common properties (both omnipotent, omniscient, etc.), they would get stuck at the step "share same origin?" – which they cannot show because Father and Son do not have the same origin (one is unbegotten, one begotten). The logic literally requires that premise, which fails. So it blocks that potential collapse. Similarly, one can't prove "divine nature = human nature" because any attempt will fail Con equality (divine nature has attributes the human nature doesn't and vice versa) and fail origin (one is uncreated, one created), etc. Good. So it keeps distinctions that should be distinct.

In summary, the readability layer not only makes it easier to work with the system by using suggestive symbols and natural reasoning patterns, but it also encodes safeguards aligned with ontological principles. This helps ensure that users of the logic (philosophers or theologians) are less likely to derive junk results – the formal language itself guides them away from trouble. It's **designed to respect the boundaries**: persons vs natures, essential vs accidental, continuous vs discontinuous, one vs many. Thus we call it *participation-ready* or *theologically safe* logic.

Having now described the user-facing aspect of the calculus, we can proceed to semantics (Section 5) to validate that all these rules make sense in actual models, and to illustrate how different "integrity policies" or "participation modes" can be set up to adapt the logic to various contexts (from purely secular to explicitly theological frameworks). We will also outline two example models: one secular for material constitution and persistence puzzles, and one theological to show Trinity and Incarnation interpretations, demonstrating that nothing in the logic contradicts those mysteries and indeed how it enforces orthodox constraints without extra ad hoc rules.

---

1 2 13 18 23 26 28 29 31 38 41 47 48 49 56 61 66 PEI-Identity-Theory.pdf
file://file_000000004e3071f584070ecee429b227

3 4 5 6 7 8 9 12 16 19 20 21 22 24 27 30 34 35 50 51 52 53 54 55 64 65 Beyond-Classical-Identity_-A-First-Order-Theory-of-Provenance-Continuity-and-Constitution (1).pdf
file://file_0000000087fc71f7a7de4cd6f5c7a393

10 CHURCH FATHERS: Council of Chalcedon (A.D. 451)
https://www.newadvent.org/fathers/3811.htm

11 62 Chalcedonian Definition - Wikipedia
https://en.wikipedia.org/wiki/Chalcedonian_Definition

14 15 17 25 32 33 44 45 46 59 60 Participatory-Relational-Logic-and-Origin-Tag-Logic_-Identity-Transformation-and-Persistence-1 (5).pdf
file://file_000000007c0071f5a22e622db92d0474

36 37 39 40 63 Part One Section Two I. The Creeds Chapter One I Believe In God The Father Article 1 I Believe In God The Father Almighty, Creator Of Heaven And Earth Paragraph 2. The Father
https://www.vatican.va/content/catechism/en/part_one/section_two/chapter_one/article_1/paragraph_2_the_father.html

[42] [43] Transubstantiation for Beginners | Catholic Answers Magazine
https://www.catholic.com/magazine/print-edition/transubstantiation-for-beginners

[57] [58] St. Thomas Aquinas: Summa Theologica - Christian Classics Ethereal Library
https://www.ccel.org/ccel/aquinas/summa.FP_Q44_A1.html