Provenance–Esse–Integrity (PEI) Framework:

A formal calculus of π-identity that extends classical identity logic

Abstract
This paper proposes a new paradigm for identity – the **Provenance–Esse–Integrity (PEI)** framework – and a formal calculus of π-identity that extends classical identity logic into hyperintensional territory while remaining conservative over it. Classical identity, symbolized "=", is preserved in all extensional contexts, but we argue that identity alone under-explains numerous puzzles. We introduce **typed judgments** distinguishing absolute identity of supposits or natures (the classical "="), true predication of a nature to a supposit ("∈"), constitution of one supposit under multiple natures ("◇"), and a **provenance-indexed identity** relation ("≡_π") that tracks origin and continuity. The guiding thesis is that identity is not a fundamental explainer but a surface verdict issued when deeper conditions of *provenance* (origin and continuity), *esse* (act-of-being), and *integrity* (sortal-specific persistence conditions) are satisfied. We formalize a logic **L_π** with new rules enforcing these conditions, proving that it conservatively extends classical logic (no pure-"=" theorem is lost or gained) while resolving classic paradoxes of coincidence, fission, constitution, and also regimenting difficult theological doctrines (Trinity, Incarnation, grace, sacraments) without ad hoc exceptions. In secular cases, the π-identity calculus cleanly separates co-located objects via constitution-not-identity[1], blocks illicit substitutions across modal contexts, and handles branching identity (fissions) by "downgrading" classical identity to the weaker ≡_π relation at the branch point. In theological applications, the same logical discipline distinguishes predication from identity, preventing the collapse of Trinitarian persons or Christ's two natures into contradiction while remaining faithful to orthodox Catholic metaphysics. The result is a single analytic framework that unifies solutions to puzzles across philosophy and theology, demonstrating that by *adding* modes of identity (rather than replacing classical identity) we gain explanatory power without sacrificing rigor or orthodoxy. However, do not be fooled, this approach may be conservative for a theological appropriation and harmony, but the theory presented here also **reframes what identity is and how it works**, using new primitives and rules that classical identity theory never formally articulated. That shift makes it **paradigm-defining**—if not fully paradigm-shattering.

The π-identity framework marks a genuine paradigm break precisely because it repositions identity from a primitive axiom to a licensed verdict. Rather than assuming identity as foundational, it treats it as an outcome dependent on deeper constraints—namely provenance, esse, and integrity—which must be satisfied before any identity judgment is even permissible. This structure internalizes substitution discipline, formally preventing cross-category substitutions that would otherwise lead to paradoxes in both philosophy and theology. Unlike systems that weaken identity or handle anomalies through ad hoc exceptions, π-identity offers a single rule-governed calculus that applies seamlessly across puzzles of fission, constitution, and theological distinction. It reframes identity not as the starting point of explanation, but as the endpoint of a prior order of continuity and

coherence. And because its explanatory power applies equally to secular and theological cases without modifying its core rules, the framework offers more than an extension—it initiates a structural rethinking of what identity is and how it functions.

The problem of substantival unity—what makes a being truly *one* rather than a mere collection or accidental heap—has long resisted reduction, sitting at the intersection of metaphysics and logic. Classical theism, especially in the Thomistic tradition, treats unity as a transcendental property of being itself: everything that exists has some mode of unity, grounded in its act of existence (esse) and informed by a unifying form. Unity is thus ontologically basic, not logically derived; it flows from the creature's participation in being. In contrast, modern identity theory, which typically begins with classical logical axioms such as reflexivity, symmetry, and transitivity, tries to capture unity through the lens of quantitative sameness or set membership. But this approach faces an explanatory shortfall: it can determine when something is the same as something else, but it cannot account for *why* a thing is one substance in the first place. For example, when confronted with a living person, classical identity theory can track continuities and labels but cannot ground the person's unity as a lived, substantial whole. The π-identity framework addresses this directly by shifting identity from primitive to derived status. Unity is no longer assumed from identity judgments; rather, identity judgments are licensed only after deeper conditions are met—namely, non-branching provenance, continuous esse, and integrity appropriate to the entity's form. This allows the system to handle unity not by stipulation or logical shortcut but through ontological coherence. Substantival unity, then, becomes visible at the point where these constraints converge, making the phenomenon of unity an intelligible effect of metaphysical order rather than a brute logical posit. In doing so, π-identity retains the mystery of being without mystifying the conditions under which something is counted as truly one, offering a framework that neither eliminates nor evades the problem but re-situates it within a layered ontology where metaphysics and logic meet.

Among the two major approaches to identity PEI builds from—the classical metaphysical tradition (Think Aristotle through Aquinas), the modern formal logic of identity (Post Leibniz)—the PEI framework emerges as the most explanatorily potent and subversively constructive. It does not discard the metaphysical insights of Aristotle, Aquinas, or Avicenna, nor does it reject the formal clarity of post-Leibnizian identity logic; instead, it reorders their priorities, placing identity as a derived judgment rather than a primitive axiom. By grounding identity in the coherence of origin, existence, and form-specific integrity, the PEI system absorbs the ontological realism of the classical tradition and couples it with the syntactic discipline of modern logic. Where classical metaphysics had depth but little formal machinery, and modern identity logic had precision but little metaphysical traction, PEI integrates both into a unified calculus that is expressive enough for theological nuance and rigorous enough for logical inference. It becomes a paradigm not by simply extending prior systems, but by formalizing the structures that earlier thinkers intuited without a framework to control their application. In doing so, it doesn't merely innovate—it completes a conceptual arc long left open. Enjoy as we develop this theory going forward.

1. Introduction and Claim

Philosophers often take the concept of identity to be simple and unproblematic: as David Lewis famously quipped, "Identity is utterly simple and unproblematic. Everything is identical to itself; nothing is ever identical to anything except itself" (Lewis 1986, 192–93)[2]. On this orthodox view, puzzles that involve identity (about change, coincidence, fission, etc.) must really be about something else – perhaps vagueness, semantics, or our conceptual scheme[3]. Yet the persistence of identity puzzles in metaphysics suggests that something deeper is at play[4][5]. This paper argues that **identity per se is not the deepest metaphysical relationship** but rather a *verdict* or outcome that becomes licensed when certain more fundamental conditions are met. We propose a new framework, **Provenance–Esse–Integrity (PEI)**, which identifies those deeper conditions and integrates them into a **π-identity calculus** (π for provenance). This calculus *conserves* classical identity in all contexts where it has traditionally been applied successfully, but it extends our expressive and inferential resources to handle cases where classical identity alone under-explains or misleads.

At the heart of our proposal is the idea that an identity statement is only unproblematically true when the items share **origin (provenance)**, have the same fundamental act-of-being (esse) in the relevant order, and maintain the **integrity** conditions of their kind over time. In everyday extensional contexts, these conditions are usually implicitly met, so classical identity "=" works as expected. But in harder cases – a statue and the clay that constitutes it, a person fissioning into two, the Persons of the Trinity sharing one nature – these conditions come apart. The π-identity calculus will explicitly track provenance, esse, and integrity to allow or forbid identity claims in a principled way, rather than forcing brute stipulations.

Claim: By introducing typed modes of predication and a provenance-indexed identity relation, we **resolve canonical metaphysical puzzles without abandoning classical logic**. We also show that the very same system can formalize key theological doctrines (like *Trinitarian consubstantiality* and *Christological dual natures*) in a coherent way that avoids both contradiction and any need for non-classical "relative identity" in the sense criticized by Geach (1967)[6][7]. The outcome is a conservative extension of classical identity theory that is *paradigm-shifting* in scope: it reframes identity as a layered concept governed by provenance (π), being (E, esse), and integrity (Int) – hence, identity as *verdict*, not *explanatory ground*.

2. Historical Accomplishment to Respect

Any new theory of identity must respect the hard-won insights of classical metaphysics and theology. In formulating PEI, we draw on a deep well of historical thought where, notably, **identity was never treated as the primary metaphysical principle**. For Aristotle,

questions of being and unity were "the hardest and most perplexing" in metaphysics[8]. Aristotle distinguished between an individual's *essence* and its *accidents*, with identity over time maintained through change in accidents but lost if the change is essential[9]. Identity was thus downstream of more fundamental categories: *ousia* (substance) for Aristotle is primary, and whether something remains "the same" depends on what kind of thing it is and what changes it undergoes[10]. Medieval thinkers built on this. **Avicenna** in the Metaphysics of *al-Šifā'* rigorously distinguished a thing's *essence* (what it is) from its *existence* (that it is), holding that all creatures are composites of essence and an act-of-being (Arabic *wujūd*, Latin *esse*)[11][12]. In creatures, essence and existence are distinct; only in God are they identical[12][13]. This implies that for creatures, *esse* (existence/actuality) is a gift or effect – in Avicenna's view, the *first effect* in the chain of causation from the Necessary Existent (God) is the existence of things[14][15]. **Aquinas** picks up this thread, calling *esse* "the actuality of all acts and perfections" (Aquinas, *De ente et essentia*; cf. *ST* I.4.1 ad3) – every created thing receives being as a primary act, before it can have any other features[16][14]. Thus identity and difference, for Aquinas, are rooted in *esse*: "Only those things can be said to be of one essence which have one being (esse)" (*ST* I.39.3 resp.)[17]. Furthermore, in classical Christian theology, *simplicity* is axiomatic for God: God has no parts, and even the divine attributes are *really identical* with the divine essence (Aquinas *ST* I.13.4). Maimonides pushed this to its apophatic conclusion – we can only say what God is not, since any positive multiplicity of attributes would compromise His absolute oneness (Guide I.51)[18]. The *indivisibility* of God carries over to God's actions: the medieval axiom *opera Trinitatis ad extra indivisa sunt* ("the works of the Trinity outwardly are undivided") insisted that Father, Son, and Spirit do not act as separate agents in creation[19]. All of these historical insights emphasize that *being* and *unity* are primary, and identity as we commonly speak of it (this is *the same* as that) is derivative. For Aristotle, *unity* is coextensive with being; for Avicenna and Aquinas, *esse* is the ground of a thing's reality, and many identities or differences trace back to whether the *act of being* is shared or not. Aquinas even says that in creatures, "no thing has existence by its essence" – it must receive existence – whereas God *is* His own essence and existence[20][14]. Thus, the tradition suggests that identity of essence (or nature) does not necessarily mean identity *simpliciter* of supposits, and identity of supposit (particular) requires a unity of origin and being.

We inherit from this history two guiding principles. First, **metaphysics precedes logic**: the logical relation of identity should not dictate our metaphysical understanding but serve it. When Aquinas says the Persons of the Trinity are "in everything one where there is no distinction of relation" (ST I.42.1 ad3)[21][22], he implies that identity holds in one respect (essence) and difference in another (relations of origin). This is not "relative identity" in the Geachian sense of an arbitrary F-relative sameness[6], but a disciplined, *typed* way of speaking about unity and distinction. Second, **esse (act of being) is ontologically prior to identity**. For created things, being is received and can be gained or lost; identity can therefore wax or wane depending on the continuity of that received being. Medieval philosophers did not formulate a modal logic of identity, but their instincts point to the need for a stratified approach: what *something is* (essence/nature), that *it is* (esse),

whence it is (its origin or cause), and *how* it persists (its integrity) are all separate questions. Identity in the narrow sense ("x = y") sits at the terminus of those questions, not at their source. The PEI paradigm thus seeks to **honor the classical achievement** – that identity is downstream of being and tightly linked to origin and kind – while systematizing it in a new logical calculus.

3. Why Classical Identity Under-Explains Hard Cases

Classical first-order logic with identity treats "=" as a primitive satisfying only reflexivity, symmetry, transitivity and Leibniz's Law (indiscernibility of identicals)[7][23]. This framework works for everyday reasoning and many scientific contexts. Yet philosophers have long noted pressure points where insisting on a simple identity/non-identity dichotomy yields paradox or mystery. We briefly review these hard cases to motivate the extensions we introduce later.

Coincidence and Constitution: Consider a clay statue (call it *David*) sculpted on Tuesday out of a lump of clay (*Lump*) bought on Monday[24]. On Tuesday, the statue and the lump coincide spatially. Are they one object or two? If one says "one", i.e. the statue = the clay, one faces the problem that the clay existed on Monday when the statue did not; thus by Leibniz's Law, they cannot be identical (one has a property – existing on Monday – that the other lacks)[1][25]. But if one says "two distinct objects", one violates the intuition (sometimes called Wiggins's truism) that two physical objects cannot occupy the exact same place at the same time[26]. Classical identity offers no middle ground between identity and absolute distinctness, yet here we have a strong sense that statue and lump are *not* simply identical (they differ in persistence conditions and kind) but also *not* like two totally unrelated objects (one constitutes the other). Philosophers like Wiggins (1968) concluded that **constitution is not identity**[1], and that the statue is "nothing over and above" the lump of clay – in other words, one object in a certain form[27]. Still, regimenting "x consists of y but x≠y" in logic is tricky: any property of x seems translatable into a property of the matter y in some form, so which differences count as genuine? This puzzle suggests we need a formal way to say "x and y share all the same matter and occupy the same location, yet they are not identical *because* they differ in form or kind." We also want to prevent illegitimate inferences like: "The lump has property P, therefore the statue has P," when P is not applicable to the statue's form (e.g., "being shapeless" might have been true of the lump on Monday, but not of the statue on Tuesday).

Fission and Duplication: Classic thought experiments (indebted to Parfit 1971) imagine a person fissioning into two – say, by transplanting each half of a brain into a different new body, yielding two successors both with continuity to the original. Call the original person *A* and the resulting persons *B* and *C*. Intuitively, B and C both stand in all the relevant psychological and physical continuity relations to A (just as a single successor in a normal case would be seen as the "same person" as A). Yet A cannot **be identical** to both B and C, because that would violate the transitivity and uniqueness of identity (if A = B and A = C, then B = C, implying only one person resulted, not two). So what is A's relationship to B and C? Many agree A is "gone" (dead, or ceased) and B, C are new, or that identity is an all-or-

nothing thing that didn't carry over. Others say identity is indeterminate in such cases. Parfit famously suggested that what matters is not identity but a continuity relation he called "R" – psychological continuity – which B and C both bear to A, and that seeing this frees us from obsession with identity (Parfit 1984). From our perspective, Parfit is pointing in the right direction: **continuity relations can survive branching even when identity does not**. We want a way to say, in our system, that B and C are each in some sense the "same being" as A *with respect to* origin and causal history, without asserting the impossible that B = C or that A survived as a single individual. In other words, there is a weaker equivalence relation – call it provenance-identity ($\equiv_\pi$) – that holds between A and B, and A and C, though not between B and C after the branch. This $\equiv_\pi$ will **not** be an equivalence relation in general (it fails transitivity in branching scenarios), which is good – it mirrors the situation. Classical identity under-describes fission: it forces a verdict of absolute non-identity, losing the important fact that "B and C *are* in one sense the same original person A." Our calculus will capture that fact in a controlled way, without violating transitivity for the strict identity "=".

Sortal-Dependence of Identity and Illicit Substitution: Philosophers like Wiggins (and Lowe, Geach, Hirsch, et al.) have argued that all identity statements are sortal-relative in the sense that one cannot meaningfully ask "a = b?" without specifying *under what concept*. A river is not the same entity as the water that constitutes it; a person is not the same as the organism or body that constitutes them (some argue). Geach (1967) went so far as to propose a theory of **relative identity** where there is no absolute "=", only "= as an F" for some sortal F[28][29]. That move is drastic and widely criticized for scrapping a fundamental logical notion[30][31]. Nonetheless, Geach's intuition was that saying "x is the same person as y" but "not the same organism as y" could be coherently true – something classical logic forbids unless one treats "person" and "organism" as shifting predicates with a single underlying identity. Similar issues arise in modal contexts: the infamous "Morning Star = Evening Star" identity is necessarily true if true at all, yet one can have true beliefs about the Morning Star that are not true of the Evening Star *as descriptions* (e.g., "the Morning Star was long thought to be a different celestial body"). In intensional contexts (belief, modality, etc.), substituting identicals can fail to preserve truth, showing Leibniz's Law has exceptions unless one carefully distinguishes use and mention, or attributes vs. names. These puzzles suggest that **unrestricted substitution of identicals is not always semantically valid** – restrictions are needed when intensionality or different "aspects" are in play. Our system addresses this by a rule of *aspect-restricted substitution*: roughly, you can only substitute terms salva veritate within the same "aspect" or mode of predication. For example, if we know x is a human being and x is also a statue (imagine a mythological metamorphosis scenario), a predicate like "...is capable of rational thought" may be valid for x under the human aspect but not under the statue aspect. Classical logic can't even express "x is a statue and a human in different respects" without collapsing them into one bundle of properties. We will introduce a notation x ◇ ⟨N1，N2⟩ to indicate x has two natures (N1 and N2) by constitution or some extraordinary process – and we will prevent the unwarranted substitution of predicates across these natures. This addresses "illicit substitution across modes."

Theological Identity Puzzles: In Christian theology, especially Catholic tradition, two central mysteries involve perplexities of identity and predication: the Trinity and the Incarnation. Classical theology insists on *Trinitarian consubstantiality* (Father, Son, Holy Spirit are "one God"/one divine essence) and *Christological unity* (the one Person Jesus Christ has two natures, divine and human). Naively formalized, Trinitarian doctrine seems to violate transitivity of identity: Father = God, Son = God, but Father ≠ Son – a logical impossibility if "=" is absolute. This led some 20th-century philosophers (e.g. Peter Geach, and later proponents like van Inwagen or Peter Brower and Michael Rea) to consider relative identity solutions – e.g., "the Father and Son are the same God but not the same Person." Our framework will show a way to avoid heretical collapse (F = G and S = G implies F = S) **without** abandoning absolute identity, by treating "x is God" as a predication of nature (x ∈ DivineNature) rather than an identity statement, and by indexing persons with their origin (provenance) to distinguish Father and Son within the one divine essence. Similarly, in Christology, we face the classic communicatio idiomatum problem: how can we say *God died on the cross* or *Jesus is omnipotent* without confusion? The council of Chalcedon says Christ is *one person in two natures without confusion,* which implies that some predicates apply to him in virtue of one nature and not the other. Standard logic doesn't allow an easy way to say "Predicate P true of x qua human but not qua divine," except via cumbersome operators or two-sorted semantics. Our calculus will allow a single subject with two channels of predication. The benefit will be a clean logical blockade against false inferences like: "Jesus is God, God is immortal, therefore Jesus is immortal (cannot die)" – which is unsound because "is immortal" must be evaluated in the channel of the nature to which it properly belongs (divine nature), and Jesus' concrete death applies only to the human channel.

In sum, the desiderata emerging from these cases are: **(a)** Conserve classical "=" for all ordinary contexts and valid inferences – we do not want to throw away standard identity or allow true contradictions. **(b)** Add *modes of being and predication* so that we can talk about one entity under different descriptions without collapsing those descriptions into one, yet without resorting to obscure metaphors. **(c)** Regiment substitution: enforce that one only substitutes identicals where appropriate to the mode (no mixing up different π-lines or natures indiscriminately). **(d)** Provide a principled criterion for diachronic identity: what makes x at t1 the same as y at t2, or not? Classical answers varied (spatiotemporal continuity, psychological continuity, etc., depending on the kind of thing). We will encode a non-arbitrary criterion via provenance and integrity: roughly, *if the causal provenance remains unbranched and the integrity conditions of the nature are not violated, identity holds; if branching or a break in integrity occurs, only ≡_π holds*. **(e)** Unify the treatment: the same core framework should handle both secular puzzles and theological doctrines, showing that theology does not require a special-case logic, only a careful extension of what is needed anyway for secular cases. In the next section, we present the PEI paradigm's primitives and judgments that meet these desiderata.

4. The PEI Paradigm: Primitives and Roles

The Provenance–Esse–Integrity paradigm introduces three key conceptual primitives and a handful of novel judgment forms:

- **Provenance (π):** Every supposit (individual entity) is associated with an origin and continuity profile, denoted by a provenance label π. Intuitively, π(x) captures *where x came from* and *how x continues to exist*. In a simple case, π could be just a unique identifier of the "causal line" of the object (like a world-line in spacetime or a genidentity relation). If an object splits or branches, we represent that as one π leading to multiple continuations (sub-π branches). **Identity in our system will require non-branching provenance**: if an object's history forks, strict identity will not hold between the branches – only the weaker provenance-equivalence (≡_π) will. Provenance thus handles cases of fission, fusion (two distinct provenance lines merging would similarly complicate identity), and any scenario where origin matters (e.g. two qualitatively identical statues made of the same clay but at different times have different provenance – they are not *numerically* identical because they didn't originate as one).

- **Esse (E):** We use E(x) to denote the *act-of-being* or first actuality of x. In creatures, this corresponds to existence – the idea that x has being by participation in existence. In God (for those sections where we apply it), E(x) for x=God is just the divine essence itself (God's essence = existence). The rule **Esse-first** in our system will ensure no created entity's identity claim can "outrank" existence: meaning, if something ceases to exist (loses its esse), all identity bets are off. This sounds obvious (dead people aren't identical to living ones, except perhaps in a loose sense like "his corpse is him in a way"), but it becomes non-trivial in theology: e.g., Christ's human nature died (lost its esse as a living human on Holy Saturday) but the Person of Christ did not cease to be – here we rely on the fact the supposit is divine and has divine esse continuous. Esse also ties into the concept of *simplicity*: in God, as classical theology holds, the persons share one identical act-of-being (the divine essence). In creatures, no two distinct individuals share the same numerically identical act-of-being – each has its own *esse* (even if the essence/universal is the same). Thus, identity of supposits will imply identity of their esse in our system. Put simply: **no created two distinct things can have the same esse** (we don't count "common being" as one numerically identical thing).

- **Integrity (Int):** For each sortal or nature F, we posit an *integrity condition* Int_F(x, t) which specifies the allowable changes in x over time such that x remains an F. This is a formal way to capture continuity of form. For example, for a human being, integrity might demand that the organism is not irreversibly biologically destroyed; for an artifact like a ship, integrity might allow plank replacement up to a threshold. **Definition (Integrity Threshold):** For a sortal F, Int_F is a condition on qualitative change such that for all times t1 < t2, if `Int_F(x, [t1..t2])` holds (x's changes from t1 to t2 stay within F's threshold) **and** π is non-branching on that interval, then

x at t1 is identical (by "=") to x at t2. If the integrity condition fails or provenance branches, then only the weaker continuity relation ≡_π may hold, not full identity. In plainer terms, integrity conditions mark the line between an essential versus accidental change for a given kind. This idea traces back to Aristotle's distinction between accidental and essential change[32], and to Wiggins's notion of sortal-dependent persistence[33]. In our formal system, we will include axioms that if Int_F holds and there's no branching, identity carries over; if Int_F fails or branching occurs, identity ceases.

With these primitives, we introduce **typed judgments** (the basic forms of statement in our logical system):

1. **Predication:** x ∈ N. We read this as "x is of nature N" or "x is an N". It means that the supposit x carries the nature or substantial form N truly. For example, if N = Human, then `Socrates ∈ Human` means Socrates truly has a human nature. In set-theoretic analogy, think of ∈ as a typing relation: it asserts x falls under type N. This is *not* a class membership in the sense of set theory, but a metaphysical predication. This distinction between saying "is a" versus "is identical to" is crucial for theology: e.g., "The Father is God" will be interpreted as `Father ∈ DivineNature`, not Father = God (since the latter would force Son = God = Father). Predication is also the vehicle for analogy: when we say "God is good" in analogical sense, we can treat that as predication with the understanding that the predicate's meaning is analogically extended.

2. **Strict Identity:** x = y. This is the classical identity of objects (supposits or of natures) **within the same category**. We only allow x = y if x and y are either both individual supposits or both natures of the same level. You cannot say "a supposit = a nature" (that's a type error, somewhat like category mistake) – for instance, it's wrong to say "Socrates = Humanity"; rather Socrates ∈ Human, and "Humanity" might be identical to "Human nature" as universals. Thus, = in our system is *sorted*: identity statements are always sortal-qualified in the sense that the context will indicate which sort they belong to. We inherit all the usual properties of identity for "=" – it's an equivalence relation and satisfies Leibniz's Law *in extensional contexts*. One key result we will show (Conservativity Lemma) is that any formula in pure "=" (with no ∈, ◇, or ≡_π) that is a theorem in our system was already a theorem in classical logic. In other words, adding our machinery does not let you prove any weird new equation between terms that you couldn't prove classically, as long as you stay in the pure identity fragment.

3. **Constitution/Aspect:** x ◇ ⟨N1, N2, …⟩. This formula expresses that **one supposit x has (or encompasses) multiple natures or aspectual parts**. We use the symbol "◇" (diamond) suggestively to indicate a composite-aspect status. For example, if we say `Jesus ◇ ⟨DivineNature, HumanNature⟩`, we mean that the supposit "Jesus" (or the Word) has two natures, divine and human. Another example: `Lump ◇ ⟨Clay, Statue⟩` might represent that a lump of clay that got sculpted now has the

form of a statue while remaining clay. In such a case, we might prefer to name the statue and the lump separately (to track provenance differences), but if one wanted to treat it as one entity under two descriptions, $\diamond$ allows that. The formal effect of x $\diamond$ ⟨N1, N2⟩ is to introduce **channels** for predication: properties or predicates can be indexed to a nature. If a predicate Q is appropriate to N1 (e.g., "can speak" appropriate to Human), then from x $\diamond$ ⟨Human, Divine⟩ one can infer Q(x) *in the Human channel*, but one is **forbidden** from automatically inferring Q(x) in the Divine channel. We state this as an **aspect-restricted substitution rule**: *From* x $\in$ N *and* N $\subseteq$ P (N is a subtype of P) *infer* P(x) (normal inheritance of predicates within a channel). *From* x $\in$ N *never infer* x = N (you cannot identify an individual with its nature). *From* x $\diamond$ ⟨N1, N2⟩ *and* "Q is true of x-in-N1 channel", we can infer Q(x) but keep it tagged as via N1; we do **not** allow substituting that into an N2 context[34][35]. In simpler terms: predications are always sortal-scoped; if x has two sortals, don't mix them. And finally, *from* x $\equiv$_π y *infer only* those predicates R which are **provenance-tracking** (for example, "has origin O", or "consists of the same matter as") but *never infer* x = y from x $\equiv$_π y. This schema encapsulates how our system avoids illicit substitution across modes and prevents conflating constitution with identity.

4. **Provenance-Identity:** x $\equiv$_π y. We call this π-identity or provenance-equivalence. It means that x and y share the same provenance/origin and continuity profile (in other words, according to the map π, they lie on the same branch or same line of development). $\equiv$_π is reflexive and symmetric, but not necessarily transitive: if y is a branch-off from x and z branch-off from y, x $\equiv$_π y and y $\equiv$_π z might hold (they share partial provenance) but x $\equiv$_π z might fail if the branch split apart. We formalize branching in terms of the π labels: think of π(x) as, say, "object #42, branch 0"; if it splits, we get π(y) = "#42, branch A" and π(z) = "#42, branch B". Then x $\equiv$_π y, x $\equiv$_π z (same root #42) but y and z are not $\equiv$_π (different branches). Provenance-identity is a weaker relation than "=", used to say objects are "the same in origin" even if they are not strictly identical. For instance, the reassembled Ship of Theseus is not strictly the same ship as the continuously repaired one in the museum (we will argue), but they are $\equiv$_π with respect to original planks/matter – one could say they are the same *material history* though not the same artifact in naval service[36][37]. We will treat $\equiv$_π as an equivalence in any non-branching context, and it downgrades to something like a similarity relation when branching occurs.

Using these concepts, we can articulate our **licensing rules** informally:

A) Esse-first: No created entity can sustain its identity through loss of *esse*. If E(x) ceases (x dies or is destroyed in the metaphysical sense), then for any y that exists later, even if we informally call it "the same" x, it's only $\equiv$_π at best, not "=" (unless God miraculously restores the same act-of-being, which would effectively be a resurrection scenario – interestingly, our system could handle resurrection as a special case of regaining the same π and a continuous enough gap maybe). This rule ensures the primacy of existence: identity is never more fundamental than existence.

B) Aspect-Restricted Substitution: As already described, you can only substitute identicals within the same aspect channel. If x = y and both are ∈ N (share the same nature), then within formulas about N or its supertypes, substitution is valid. But if x has two different natures in play, you cannot use an identity or predication from one nature's context to another directly. For theologians, this blocks the fallacy "Jesus is human (true), Jesus is God (true), therefore some human is God (invalid cross-substitution)" – the last inference tries to equate the nature "human" with the supposit "Jesus" and substitute into the "God" context, which our typing forbids.

C) Integrity and Non-branching (Diachronic Identity): If from time t0 to t1, object x's nature integrity holds (no essential change) and π has no branching, then we license x(t0) = x(t1). If either an essential change occurs (integrity fails) or a branching (two continuers), then we stop short: you may have x(t0) ≡_π y(t1) for each continuer y, but not x = y. A special case: if integrity fails drastically, x might cease to be ∈ N and become ∈ N' (a different nature) – that is effectively a new identity (like transubstantiation in theology or sci-fi metamorphosis: the original is gone in that form, a new entity with new form is present sharing maybe matter or something).

D) Constitution-not-Identity: If x and y are co-located and one constitutes the other (say x ◇ ⟨F,G⟩ where G is the material or lesser kind and F the higher-order form), then x ≠ y, but we may have x ≡_π y if historically they share origin. For example, Lump and David (statue) earlier: Lump (clay) becomes constituted as David (statue). We represent David ◇ ⟨Statue, Clay⟩ or David ∈ Statue and David's matter ∈ Clay. In either case, David ≠ Lump (their identity is kept distinct because one has two sortals; essentially, David is Lump-under-a-form, but we don't identify them to avoid Leibniz Law issues[1]). Instead, we assert a constitution relation (which could be encoded via ◇ or a separate predicate Con(x,y)). The rule is: never conclude x = y just because x constitutes y or shares all the same parts. This directly counters the argument that "no two things can be in the same place; so if they are, they must be one thing." We bite the bullet that two "things" (in different senses) can coincide – and our logic can handle it by differentiating their modal properties via the sortals and integrity constraints (the statue can survive smashing if reassembled, the lump cannot survive being reshaped into a bust, etc., or vice versa depending on definitions).

E) Branching Downgrades "=" to "≡_π": This we have covered – at a branch point, what would have been one continuing object becomes two provenance-connected objects. We replace identity claims with provenance-equivalence claims. This ensures classical identity's transitivity is preserved globally: we simply don't assert a = b in those cases.

Overall, the PEI paradigm doesn't eliminate the classical concept of identity; it **adds layers** around it so that when we utter x = y, it is only after passing checks of same provenance (no branching), continuity of being, and satisfaction of sortal integrity. In everyday cases, these checks are trivially passed, so the extra apparatus stays in the background. In hard cases, the apparatus activates and guides us to say "not equal, but something else holds between them."

The next section formalizes these ideas into a deductive system L_π.

5. Formal System L_π: Syntax and Inference Rules

We now present **L_π**, a formal calculus embodying the PEI paradigm. Our approach is natural deduction style (Fitch-style) for readability, with clear introduction and elimination rules for each of the new logical symbols (∈, ◇, ≡_π) as well as modified identity rules.

Language of L_π:

- **Terms:** We assume distinct sorts of terms for *supposits* (individual entities, like objects or persons) and for *natures* (kind terms). We will use lowercase Latin letters a, b, c,… for supposit constants, uppercase letters like N, F, G for nature constants, and possibly indexed terms for time (e.g., a@t for "a at time t") or we introduce time as a parameter in formulas when needed. We include a special set of *provenance labels*; however, we won't usually write π labels explicitly in terms, rather π is an implicit function mapping each supposit term to a label. For formal semantics, we might treat the domain of discourse as pairs (entity, label) but in syntax, we keep it cleaner.

- **Atomic Formulas:** These include t1 = t2 (where t1,t2 are terms of the *same sort* – either both supposits or both natures), t1 ∈ N (t1 a supposit term, N a nature term), t ◇ ⟨N1,...,Nk⟩ (t a supposit, N_i nature terms, meaning t has all those natures; typically k=2 in our use cases but one could have many), and t1 ≡_π t2 (supposit terms t1, t2). We also allow standard predicate symbols which may be indexed with a nature if needed. For example, we could have a unary predicate Mortal(x:Human) indicating x is mortal *in virtue of being human*; but to avoid complicating notation, we often can treat such predicates as ordinary and rely on the context of ∈ to know where it applies.

- **Classical connectives:** ¬, ∧, ∨, →, ∀, ∃ all included, with usual formation rules.

- **Logical Axioms (Identity Axioms):** We include the classic axioms: ∀x (x = x) (reflexivity); ∀x∀y (x = y → y = x) (symmetry, or derivable); ∀x∀y∀z ((x=y ∧ y=z) → x=z) (transitivity); and Leibniz's Law restricted: if φ is an *extensional* formula (one not containing ∈, ◇, ≡_π or any intensional context like modal operators), then from x = y we can infer φ(x) → φ(y). We *do not* allow substitution of equals blindly in formulas with the new operators or in intensional contexts without checking conditions (those will require additional rules below).

Introduction and Elimination Rules:

- **Predication (∈):** *Introduction:* If term t is declared of type (or nature) N in the context (we assume each constant might come with a declared nature in the model), we can introduce t ∈ N. More interestingly, we have rules to infer larger sortal membership from smaller: if N ⊆ P (as nature types, N is a subtype of P) and x ∈ N, then infer x ∈ P. (For example, if Human ⊆ Animal and Socrates ∈ Human, infer Socrates ∈ Animal.) *Elimination:* From x ∈ N one cannot directly derive much except

that x is available for nature-specific rules (for instance, if an axiom says "All Fs have property G", and we have x ∈ F, we can infer G(x)). We caution again: one never treats ∈ as set membership in a comprehension sense, so there's no "extensionality" or so for ∈.

- **Constitution (◇):** *Introduction:* If we have x ∈ N1, x ∈ N2, ... x ∈ Nk and certain side conditions (for example, the N's are pairwise compatible or come from a known construction), we may introduce x ◇ ⟨N1,...,Nk⟩. Typically, we might have an axiom schema: if N1 and N2 are not logically incompatible (e.g., Divine and Human are incompatible in the sense no *natural* being can be both, but in the case of Christ we allow it exceptionally), then one can have a single supposit with both. In general, introduction of ◇ might come from a rule that says "there exists an x such that x ◇ ⟨N1,N2⟩" if certain premises are satisfied (like one nature being assumed by another). *Elimination (aspect switching):* From x ◇ ⟨N1,N2⟩, we cannot eliminate to get any simpler formula directly, but we use it with specialized rules: if we know x ◇ ⟨N1,N2⟩ and we have a predicate or theorem φ that applies under aspect N1, we can assert φ(x) with a tag [via N1]. We might formalize this by indexed predicates: e.g., from N1(x) we get N1@N1(x) (trivial) and from that plus x ◇ ⟨N1,N2⟩ we infer just N1(x) (since x indeed has that property as an N1). However, we do not allow concluding N1(x) implies N1@N2(x) or something silly. Essentially, ◇ ensures that x can be treated as an N1-thing *and* as an N2-thing, but one must keep track of which way one is viewing x in any sub-proof. In practice, it might be simpler: we treat statements like "P(x)" as shorthand for "P_N(x)" meaning P is being said of x under nature N, which is disambiguated by context. Then x ◇ ⟨N1,N2⟩ allows both P_N1(x) and Q_N2(x) to be true for different P,Q, and it blocks mixing them up.

- **Provenance-Identity (≡_π):** *Introduction:* We introduce x ≡_π y when we know x and y share the same provenance label and time sequence and integrity up to a point. One inference rule: if x = y then certainly x ≡_π y (identity implies provenance-identity). Another: if π(x) = π(y) (in the meta-language, labels identical) but perhaps x ≠ y in object-language because we treat them as distinct nodes, we can introduce x ≡_π y. For example, after a fission, we might treat the original and the copy as different objects in semantics but they have a common origin label, so in logic we allow ≡_π. *Elimination:* The main use is that from x ≡_π y, one can infer that any predicate that is explicitly marked as provenance-insensitive or continuity-based will be shared. For instance, if R is "was created in Athens in 500 BC" or "has serial number 42", those are provenance facts, so x ≡_π y implies R(x) ↔ R(y). But ≡_π elimination does **not** allow substituting y for x in arbitrary formulas; in particular, it doesn't allow concluding x = y. So ≡_π elimination is very limited – it's mostly there to carry information about origin and perhaps to chain: if x ≡_π y and y ≡_π z *and no branching occurred*, one might infer x ≡_π z. But if branching occurred, we explicitly will have a premise that says "if π is branching between y and z, then you cannot infer x ≡_π z except perhaps a weaker statement." So the logic will have a *Non-Branching Identity rule*:

Non-Branching Identity Rule: If (i) x ≡_π y, (ii) the provenance π is non-branching for the interval covering x and y (meaning x and y lie on a single continuous line of descent with no splits), and (iii) for all times t0,t1 from the existence of x at t0 to y at t1 the integrity conditions Int_N(x, [t0..t1]) hold for the relevant nature N of x, *then* infer x = y. In natural deduction, this might be an introduction rule for "=" (showing x and y are identical by criteria). Contrapositively, if x ≠ y (negation), then either branching occurred or some integrity was violated, or they had different esse at some point.

- **Identity (=):** *Introduction:* Typically by the rule just above (non-branching plus integrity yields identity), or trivial ones like if we are within a subproof assuming x and we deduce x, then by reflexivity we conclude x = x (but that's usually given axiomatically). We do not generally allow arbitrary introduction of x = y unless we have proved all their differences vanish – essentially we treat identity as strong, only conclude it when absolutely warranted by previous ≡_π and no differences. *Elimination (Substitution):* We restrict Leibniz's Law in a **mode-aware** way. In formal terms: if φ is a formula and we know x = y, we can replace x with y in φ **only if** φ is a formula in the extensional fragment or if every occurrence of x in φ is under the same aspect or context. If x appears in multiple sorted roles in φ, substituting y might cross channels which is not allowed. To implement this, one approach is to stratify the language into layers and only allow substitution in the bottom layer (classical part). Another approach: treat ∈, ◇, ≡_π as "non-logical symbols" with extra side conditions and simply say we won't substitute inside those atomic formulas – instead, we use their own rules. For clarity, we can state: From x = y and x ∈ N, infer y ∈ N (if sorts match). From x = y and x ◇ ⟨...⟩, infer y ◇ ⟨...⟩ (since if x was the supposit with multiple natures, y is the same supposit). Likewise x = y and x ≡_π z gives y ≡_π z, etc. Those are legitimate because identity just allows renaming the subject. But we won't allow going from x = y and (x ∈ N and x ∈ M) to infer y ∈ N and x ∈ M then something mixing y and x. Actually if x=y, you could trivially drop x or y anyway. The main thing is we do not allow: x = N and x = y to infer y = N (since x = N is type error in our sorted approach to begin with).

To illustrate the use of these rules, we provide a few **sample derivations** (worked in Fitch-style outline):

(i) Co-location without Identity (Statue and Lump): We assume constants *d* for the statue "David" and *l* for the lump "Lump". Premises: d ∈ Statue, l ∈ Clay, and at time t1 just after sculpting, we have d ◇ ⟨Statue, Clay⟩ and l ◇ ⟨Clay, Statue⟩ (meaning the lump l now has the form of a statue; possibly we treat d and l as the same physical matter but conceptually two aspect-entities). We also have provenance labels: suppose π(d) = π(l) = some label (they originated together when the sculptor formed it). Now, we want to show ¬(d = l). Assume towards contradiction d = l. Then d and l have exactly the same properties. But consider property P = "was an unformed lump on Monday". l had P true (since l existed as just clay pre-sculpture), d did not (d as a statue did not exist on Monday). By Leibniz's Law in extensional context: P(l) and ¬P(d) would violate d = l. So contradiction. Therefore ¬(d = l). However, we can assert d ≡_π l because they share origin. And we have

both d and l co-constituting the same material at t1. In L_π, one formal way: from d ◇ ⟨Statue, Clay⟩ we know "MaterialClay(d)" and from l ◇ ⟨Clay, Statue⟩ we know "FormStatue(l)". These are not contradictory because they are on different channels. We do get that the weight of d = weight of l, etc., by ≡_π or by the fact they share all parts. But any inference like "d has shape because statue, thus l has shape because l = d" is blocked since l is considered under Clay aspect which might not guarantee the same persistence of shape (clay can be reshaped). So the rules successfully block confusion. The result is a **consistent description**: statue and clay coincide but are not identical, yet no double-counting of matter or violation of physical facts occurs because one is constituted by the other[38].

(ii) Branching (Fission case): Let A be original person, B and C the two offshoots. At initial time t0, A exists with π = α. At t1 post-fission, B has π = α_A, C has π = α_B (two branches off α). We can derive A ≠ B by non-branching rule contrapositive (since branch occurred, identity cannot be maintained). Similarly A ≠ C and obviously B ≠ C (distinct people). But we assert B ≡_π A and C ≡_π A (common origin α). Now suppose someone claims "transitivity of identity says B should equal C via A" – but we never said A = B or A = C, so classical transitivity is not violated. We only have the ≡ relation which is not transitive here. Thus no contradiction. Yet we have meaningful statements: B and C are in one sense the same person as A was (say, same memories, etc.), so we can express "B and C are continuation of A" as B ≡_π A ∧ C ≡_π A. We might also define a predicate Alive(x,t) and prove that if Alive(A,t0) and after fission Alive(B,t1) and Alive(C,t1), then "A is alive at t1" is false strictly (since A ended), but something of A survives in B and C. Parfit's point about caring for continuity rather than identity can be captured: any predicate like "has A's personality" could be proven true of B,C by provenance equivalence and integrity of psychology perhaps, without asserting identity.

(iii) Substitution discipline example: Suppose Jesus is represented by constant J. We have J ◇ ⟨Divine, Human⟩. Thus J ∈ DivineNature and J ∈ HumanNature. Let's formalize a problematic inference in naive logic: "God is immortal. Jesus is God. Therefore Jesus is immortal." In L_π terms: "God is immortal" – we interpret "God" here as a nature-term or perhaps as the property "∀x (x ∈ Divine → Immortal(x))". That is a doctrinal truth: any supposit with divine nature is immortal *in respect of the divine nature*. "Jesus is God" we interpret as J ∈ DivineNature. From that and the axiom, we can derive Immortal(J) *in the divine aspect*. But the conclusion "Jesus is immortal" usually means "immortal in his whole person" or specifically we might misapply it to the human nature. Our logic would say: yes, Immortal(J) via Divine. But we also have Mortal(J) via Human (because humans are mortal, Jesus had a passible nature and indeed died). So we end up with two statements: Immortal^Div(J) and Mortal^Hum(J). These are not contradictory because they are indexed by aspect. If one forgot the aspect and treated them as one, one would get a direct contradiction. The system's syntax prevents that by not allowing those two predicates to be used in the same context without distinguishing their aspect. In a meta-sense, one could say "In one respect Jesus is immortal, in another he is mortal" which is exactly the orthodox view. We can even prove that you cannot infer a contradiction:

assume for contradiction that from the premises you got both "Immortal(J)" and "¬Immortal(J)". In our logic, to derive ¬Immortal(J) you'd need to derive Mortal(J) and equate that with ¬Immortal (given mortal means not immortal). But Mortal(J) would come from J ∈ Human and rule that humans die. Immortal(J) came from J ∈ Divine and rule that God cannot die. These two cannot be combined unless we drop the distinction; any rule that tries would violate aspect-restricted substitution, hence it's not derivable formally. Thus, no contradiction arises, solving the communicatio idiomatum logic puzzle. Similarly, for Trinity: "Father ∈ Divine, Son ∈ Divine, Father ≠ Son" is consistent, and we can show that from that one cannot derive "Father = Son". A naive derivation would have treated "Father is God" and "Son is God" as identities "Father = God, Son = God" and then by transitivity Father = Son, contradiction. In our system "is God" is ∈ `Divine`, not an identity, so transitivity of identity isn't triggered at all. We will provide a consistency proof sketch in Section 8.1.

Finally, we note that L_π can be formulated as a sequent calculus or implemented in a proof assistant. It is essentially a multi-sorted first-order logic with some extra predicates and a careful restriction on substitution. We ensure **every theorem in pure identity logic holds** (reflexivity, symmetry, transitivity of = remain, etc., since we didn't mess with them except restrict where we apply Leibniz's Law which is anyway outside the pure identity tautologies). We also expect usual **consistency** and **soundness** to hold; in Appendix A we outline a formal semantics to verify this.

6. Semantics and Metatheory

To prove the rigor of L_π, we define a formal semantics. Models will show how the new symbols are interpreted and allow us to prove soundness and completeness relative to these models (a full proof theory is beyond scope, but we sketch key results).

Model Structure: An L_π model $\mathcal{M}$ consists of:

- A nonempty domain D of *supposits* (individuals).
- A set N of *natures* (which might be seen as intensional entities or simply labels for subsets of D).
- An interpretation function for constant symbols of each sort (mapping each supposit constant to an element of D , each nature constant to an element of N).
- A function $E: D \to$ Truth or a subset $E \subseteq D$ indicating which supposits "exist" or rather, since we consider possibly temporal existence, we may have $E(x, t)$ meaning x exists at time t. But we could also incorporate time by having domain vary over time; a simpler approach is to treat time as an index: our satisfaction relation will be relativized to a time parameter $t \in T$, where T is a linearly ordered set of times.
- A **provenance function** $\pi: D \to P$ where P is a set of provenance labels structured possibly as a rooted tree. Each label π(x) indicates the origin ID of x (like an equivalence class for origin) possibly with a branch ID.

- An **integrity function** that for each nature $F \in N$ provides a criterion for persistence: formally, we can have a function $\text{Int}_F : D \times T \times T \rightarrow \{\text{true,false}\}$ such that $\text{Int}_F(x, t_0, t_1)$ is true if and only if "x retains nature F integrity throughout interval $[t_0, t_1]$ ". In practice, this is model-specific but it must obey certain properties (transitive: if it holds from t0 to t1 and t1 to t2, then from t0 to t2, etc.).

- A **constitution relation** $C \subseteq D \times N$ that interprets $x \in F$. We may also include an explicit subset $A \subseteq D \times \mathcal{P}(N)$ for the diamond operator: $(x, \{N_1, \ldots, N_k\}) \in A$ iff $x \Diamond \langle N_1, \ldots, N_k \rangle$ is considered true. However, we can derive that from C: basically $x \Diamond \langle N_1, \ldots, N_k \rangle$ holds iff $x \in N_i$ for all i (and perhaps maximal such set). So we might not need A explicitly, and simply define the truth of $\Diamond$ in terms of multiple membership in C.

Truth Conditions (at a time t):

- $\mathcal{M}, t \vDash x \in F$ iff $\big(val(x), val(F)\big) \in C$; i.e., the supposit denoted by x is of nature F in the model. If our domain is typed, we ensure val(x) is an element of val(F) if we think of val(F) as a subset of D. Another approach: each nature denotes a set of possible supposits (the extension of that nature at time t). For dynamic nature (someone might cease to be human? usually not, natures typically stick until you die, except in weird metamorphoses or resurrection contexts). We can consider that if x is alive or extant then x ∈ Human implies x's current state qualifies as human. We won't complicate further – treat ∈ as basic predicate relation.

- $\mathcal{M}, t \vDash x = y$ iff $val(x) = val(y)$ as elements of D. This is standard.

- $\mathcal{M}, t \vDash x \equiv_\pi y$ iff either $val(x) = val(y)$ (trivially same origin if same object) or else $\pi\big(val(x)\big) = \pi\big(val(y)\big)$. That is, x and y have the same provenance root. We might refine: require also that up to time t, their line hasn't branched apart. But that branching info can be encoded in π's value including branch info. For example, we can say π returns a pair (rootID, branchID). Then x ≡_π y means rootID is same. Non-branching can't be evaluated in an instant semantics easily without some history. So for semantics, we might cheat: define a binary relation R on D such that R(a,b) is true iff a and b share all historical stages up to time min(time(a), time(b)). This is complicated; simpler: we give semantics to formulas over intervals. But to keep it simpler: we interpret ≡_π just as same origin label.

- $\mathcal{M}, t \vDash x \Diamond \langle N_1, \ldots, N_k \rangle$ iff for all i, $\mathcal{M}, t \vDash x \in N_i$. And typically, that set {N1..Nk} is maximal such that this holds (no bigger set of natures all true of x). But we don't need maximality in truth condition, any superset wouldn't be used in syntax presumably.

- For integrity and non-branching: these are meta-conditions rather than atomic formulas in our logic (we didn't include an operator to say "π is branching" – though we could have a predicate Branch(π,x,t)). Instead, we incorporate them in inference rules. In the model, we say: identity over time is real identity in the model. If at time t0 and t1 we refer to the same object val(x) and val(y), identity's truth conditions require they be same element and time points maybe irrelevant. We want: if x at t0

and y at t1 are meant to be same object, model must give them same element. And if not, then not. We ensure that by how we interpret identity.

Key Results:

- **Soundness:** Every inference rule of L_π preserves truth in all models. For example, our restriction of substitutivity is such that if x = y in a model, any extensional context sees them interchangeably. In intensional contexts, we didn't allow substitution, so we preserve truth vacuously. The more interesting ones: Non-Branching Identity Rule – if π(x) = π(y) and integrity holds from t0 to t1, then model should have given x and y same element. We design models such that an object under continuous conditions doesn't split. So presumably each provenance root corresponds to an equivalence class of possibly many elements over time (like temporal parts maybe), but if integrity holds, we ensure only one element covers that span. If we had multiple, that'd violate uniqueness. So in any model, if the premises are satisfied, it must be that val(x) = val(y), hence x = y is true, making rule sound. We won't formalize fully here due to space.

- **Conservativity Lemma:** If a formula φ in the language of ordinary first-order logic with "=" only (no ∈, ◇, ≡_π) is provable in L_π, then φ is logically valid in classical logic. Conversely, if φ is a classical logical truth, it's provable in L_π (since we haven't removed any classical axiom, just added more). This is evident because any proof in L_π of a formula without the new symbols couldn't use the special rules except trivial identity ones which coincide with classical rules in such a case. We thus haven't changed classical identity's theory at all; we've only added the ability to speak and deduce extra things.

- **Substitution Discipline Theorem:** Roughly, this metatheorem says: If a formula is derivable that improperly mixes modes (like it concludes something that would correspond to an invalid cross-context substitution in naive logic), then there must have been a step that violated our rules. In other words, our calculus guarantees that no proof can equate or confuse an entity's distinct aspects or identities. Formally: There is no derivation in L_π of a formula of the form "x = N" (supposit equals a nature) or "x = y" where x, y have different declared sorts or provenance (no proving a clay = a statue in same sense), etc., unless the premises themselves were inconsistent. This is a kind of type-safety result: terms of different category cannot be proven identical. Similarly, one can prove that if x ≡_π y is derivable, either x and y came from a common identity chain or the theory has a contradiction if one tries to push it further to equate them despite a branch.

- **Branching Lemma:** If a provenance label π branches at time t (meaning there exist distinct objects y, z such that y ≡_π x and z ≡_π x and at times >t they are not identical), then there is no object w at time > t such that w = x (the original). In other words, identity stops at the branch; beyond it we can't find a single identical successor of x. This lemma is reflected in the logic because we would not be able to

prove x = y or x = z after branching (only ≡). In the model, this corresponds to how we interpret branching: either treat it as actual splitting of an object into two new ones (which in many metaphysical interpretations is impossible for concrete objects unless we accept multi-location or 4D worm views). But in any case, the logic enforces that you cannot pretend they remained one.

Thus, the metatheory assures us L_π is a **conservative, consistent extension** of classical FOL=, with extra expressive power and constraints that mirror intuitive metaphysical principles.

7. Secular Flagship Puzzles Resolved

In this section, we apply the π-calculus to several notorious puzzles of identity, showing how each is resolved or dissolved. The examples will illustrate the use of our new notations and rules "in action."

7.1 Statue and Lump (Material Constitution)

The Puzzle: A sculptor molds a lump of clay (call it *Lumpl*, L for short) into a statue named *Goliath* (G)[24]. At the end of sculpting, every particle of clay in Lumpl composes Goliath. Yet several intuitive differences hold: Lumpl could survive squashing (it would just be a shapeless lump again), whereas Goliath would not (a squashed statue no longer counts as that statue). Lumpl existed a day before, Goliath did not[25]. So Lumpl and Goliath have different properties, hence by Leibniz's Law they are not identical. But at that moment, they occupy the same space and are made of the same matter – so how can they be two things? This is the classic "coincidence" or material constitution problem[1].

PEI Resolution: We say: Goliath and Lumpl are **two designations with two natures for one continuous hunk of matter**. In our framework, we can formalize the situation as follows. There is one underlying supposit, call it *x*, such that initially (before sculpting) x ∈ Clay and after sculpting x ∈ Clay and also x ∈ Statue. We could model this as one supposit with changing nature (which violates integrity of the statue form because it didn't have it earlier). But a clearer way is to actually treat them as distinct supposits that coincide – effectively, Lumpl (l) and Goliath (g) as separate supposit terms with a constitution relation. So at t1 (after sculpting): g ∈ Statue, l ∈ Clay, and g ◇ ⟨Statue, Clay⟩ and l ◇ ⟨Clay, Statue⟩. The meaning of those last statements: g is a statue constituted of clay at t1, l is a lump constituting a statue at t1. How can we have two supposits share the same matter? In the model, we might say they share all proper parts; or we consider that g and l are actually the same *hunk of matter* but conceptualized differently. If we treat them as one supposit, then it has two natures, so better: let's consider them one entity with two names. Actually, to avoid confusion, let's do it with one supposit x: x ◇ ⟨Statue, Clay⟩. Then Goliath and Lumpl are just two names for x in different roles. That is arguably the *one-object* interpretation (constitution as identity, but under different sortals) which is close to a relative identity or sortal essentialist approach (Wiggins's approach was similar but insisted they're not identical tout court[27]). Alternatively, we maintain they are distinct but

coincident objects: then how to avoid them being identical? The answer: they have different π or E or Int that prevents identification. Indeed, we can say: they share all current matter and location, but differ in **provenance**. Lumpl (l) existed since the clay's creation, say π(l) = "ClayBatch#5 created Monday". Goliath (g) came into being when sculpted: π(g) = "SculptureEvent#7 on Tuesday". So g and l have different provenance labels – that's a crisp difference. Thus, even though from Tuesday onwards they move and change together, their origins differ, so we cannot say g = l because that would entail they were one being all along (which they weren't). In our logic: g ≠ l because, for instance, g has the property "being created on Tuesday" which l lacks (l was created earlier)[25]. However, g ≡_π l might be false here since we gave them different π – so they're not even provenance-identical. If that feels counterintuitive (should they share a provenance from molding onward?), we might refine provenance to allow temporary convergence – but that's too fine. Simpler: say provenances are different, so no ≡_π either. Then what relates them? The constitution relation: g ◇ ⟨Statue, Clay⟩ and l ◇ ⟨Clay, Statue⟩ in effect tie them. We also impose an **integrity threshold**: For the sortal Statue, Int_Statue(x, t1..t2) demands no severe shape change; for sortal Clay, Int_Clay(x, t1..t2) demands maybe conservation of mass and cohesion. At t2, if we smash the statue, what happens? The clay persists (Int_Clay satisfied, no branching; so l at t1 = l at t2 by rule), but the statue does not (Int_Statue fails at smashing moment, so g at t1 has no identical g at t2; the statue ceased). Our logic thus yields: Before smashing, we could deduce l = l (trivial), and perhaps l ≠ g always. After smashing, l still exists, g does not. No contradiction. And importantly, no double-counting: any statement about weight or spatial occupation – we formalize those as predicates that likely ignore sortal, so either we say "Mass(x) = 5kg" was true of x the clay-statue; it doesn't imply two 5kg objects coincide giving 10kg; it's one mass predication true under either description. If we had them as two objects, we'd need to say the mass of statue is a derivative property of the clay's mass. Our system can accommodate either approach.

To sum up: **Why S ≠ L?** Because they differ in provenance and essential properties. **Why co-location doesn't force identity?** Because our logic allows one supposit to carry two natures or two coincident supposits to share matter via constitution rather than identity. It blocks the inference that co-location → identity by explicitly representing the alternative (constitution).

In philosophical terms, we provide formal underpinning for the intuition some have expressed: "The statue and the clay are the same material object in one sense, but not the same artifact." Our system just makes the senses precise: "same in material provenance (≡_π_matter), not same by strict identity." In fact, we can create distinct provenance types: π_matter vs π_form perhaps, to say they coincide materially but not formally. But that sophistication isn't needed here.

This resolves the statue-lump paradox cleanly: no need for mysterious metaphysics like temporal parts or brutal essentialism beyond our formal Int_F constraints – we codify the essence vs accident difference, and the logic handles the rest. Notably, **Leibniz's Law is preserved** but only applied in proper contexts. The property "existed on Monday" is sortal-

independent, so from l and g being the same at Tuesday, one might try to propagate that property – but since we never accepted g = l, there's no substitution to do. And we don't want g = l precisely because of that differing property. So everything is consistent.

7.2 Fission and Fusion (Persons and Organisms)

The Puzzle: Imagine a science-fiction scenario of teleportation gone wrong that causes one person to split into two, as in Parfit's famous example. A enters the teleporter, and two identical copies, B and C, emerge. Who is A? Intuition: A psychologically survives as both B and C (each believes they are A), yet obviously B and C are distinct people after. This undermines the notion that identity is what matters in survival – since identity cannot be one-many, A can't be identical to both, so perhaps A ceased and two new persons came (which seems to conflict with the continuous experience of at least one, perhaps we say the "original" experience bifurcated). Fusion is the reverse: think of two amoebae merging into one organism – two origins, one outcome; which one survived? Maybe neither, a new one.

PEI Resolution: Non-branching or no identity is our slogan here. According to Rule (E) earlier, when provenance branches, strict identity stops[39]. So in A's case: Let's give A label π = α. When teleporter duplicates, we assign B label α_1 and C label α_2, indicating they share root α but diverged. Now, in L_π we will not assert B = A or C = A, because doing so would violate the one-one nature of "=". Instead, we assert B ≡_π A and C ≡_π A (common origin) and also that from A's last moment to B's first, the integrity conditions (psychological continuity, say) hold along branch 1, similarly for C along branch 2. The **Branching Lemma** we proved guarantees no single identity continuation exists; indeed, our Non-Branching Identity rule's contrapositive says: if you had A = B, then provenance α must not have branched – but it did, so A ≠ B. Likewise A ≠ C. And obviously B ≠ C since they're two people standing apart with different π branches.

Our system thus predicts: post-fission, A is gone in the sense of numerical identity; two new persons exist, each *provenance-identical* to A (meaning each carries A's history as their own pre-split history)[4][40]. This fits the idea some have that fission is death of the original with two offshoot "children" that share memories. Others might prefer to say A survives "in" B and C without being identical – our ≡_π captures that notion formally. It's akin to Lewis's idea of person-stages: A had two future continuants but no single continuation. The π calculus can be seen as a more fine-grained variant of that, where π tracks the tree of continuants.

One might worry: is transitivity of identity broken? Not at all: A = A, B = B, etc. Transitivity just isn't applicable because we never assert A = B or A = C. Provenance-equivalence (≡_π) is not transitive in this scenario: we have B ≡_π A and A ≡_π C, but B not ≡_π C due to the branch difference. So we cleanly represent the non-transitive "survival" Parfit talked about[41].

The system also blocks any unwarranted inference like "B and C are both A's continuers, so B and C are the same person" – that would require B ≡_π C and some assumption like "if

two things are ≡ to a third, they are ≡ to each other." That property holds only if no branching, but here branching, so it fails. We might formalize: if B ≡_π A and C ≡_π A and not Branch(α) then B ≡_π C and then by non-branching identity maybe B = C, but Branch(α) is true here, so the antecedent fails, no contradiction.

For fusion: Suppose two amoebae X and Y merge into Z. Then π(Z) might have two roots (if we allow that), or we say pick one as main and note that Y's lineage ended by assimilation. Either way, ≡_π would not apply straightforwardly (since two origins yield one result). This could be seen as the inverse of fission and just as problematic for classical identity (backward branching in time). Our system similarly would not allow both X = Z and Y = Z, since that'd imply X = Y. Instead, one or both are gone, with Z a new or continuation of maybe X (if Y is absorbed). One could say Z ≡_π X and also Z ≡_π Y but that violates how we set up ≡ (we could allow an object to have multiple provenance though—maybe not, probably better to treat one as main). We won't digress; suffice that such cases can be handled by careful assignment of provenance and acknowledging identity failure for one line.

Thus, the π-identity calculus resolves fission by refusing to force a binary identity judgement, offering instead a structured continuity relation. It "dissolves" the puzzle in line with Parfit's view: yes, something is lost (the strict identity of A), but something important (continuity of person) remains, now distributed between B and C[42]. And our logic can even accommodate normative or psychological predicates to formalize Parfit's claim that what matters (survival, memory continuity) can be had without identity.

7.3 Ship of Theseus (Continuity vs. Original Parts)

The Puzzle: A famous ship has its planks replaced one by one over years until none of the original wood remains; is it the same ship? Many say yes, due to continuous use and structure. But suppose all the old planks were stored and eventually reassembled into a ship – now which is the "real" Ship of Theseus[43][36]? We end with two ships both claiming identity with the original: one by continuity of form, one by reunion of matter. It seems we can't grant identity to both without contradiction, yet each has a plausible claim[44].

PEI Resolution: The PEI approach neatly sorts out the claims by distinguishing provenance types and integrity conditions. We define two kinds of provenance in this story: one for *form-continuity* (the ongoing organization in Athens), and one for *matter-continuity* (the planks). Initially, the Ship of Theseus (call it S) has provenance label π = ("Theseus's ship" structural line, id #Th) and also each plank has a material provenance id. As parts get replaced, the ship S keeps the same structural π (#Th stays with it, non-branching) because the crew maintained the ship's identity through repairs – integrity Int_ShipStructure was preserved: each small change was within tolerance (the ship's form changed gradually, always sea-worthy, etc.). Meanwhile, the matter that left – each plank – went into a warehouse. Eventually, those planks are assembled into a ship S'. Now, do we call S' the same as S? In our terms: S (in museum now, fully replaced planks) has

provenance #Th, branch A (no branching actually on structural line, it's single). S' (reassembled) gets a new structural provenance, say #Th_matter. But its material components have provenance connecting to original planks.

We articulate: **Non-branching structural provenance is the criterion for identity of artifact.** So the ship in Athens that underwent gradual repair qualifies as *the* Ship of Theseus – let's say S_original = S_continuing (it never stopped being Theseus's ship). The reassembled one we'll call S_reassembled. Our logic says: at final time, S_continuing = S_original (trivially), S_reassembled ≠ S_original (they are distinct ships now, coexisting). But what about S_reassembled's relation? We say S_reassembled ≡_π^matter S_original (they share origin in terms of material). Perhaps define ≡_π_m to denote same set of original parts. Meanwhile, S_continuing ≡_π^form S_original by form continuity (actually = since no branching on form line). If we had to pick "the" Ship of Theseus, we favor the one with form continuity as truly identical, and treat the other as a successor in matter only, thus not the same ship but a historically continuous object of another sort (like a replica but made of original parts).

This matches one common resolution: The original ship is the one in the museum; the plank-ship is a different ship (sometimes jokingly called "Theseus's ship's ship" or just a replica). Our system defends that by noting the branching: when the custodian started collecting planks and planning reassembly, a second potential ship identity started branching off. But that branch was not realized as an actual ship until all parts assembled later. By then, the original's structural lineage had long diverged. So at the time of reassembly, you truly have two separate ships with a common origin story. And indeed in our logic we'd not allow "S_reassembled = S_original" because that would conflict with the fact that S_original had continued existence elsewhere concurrently (you can't have one thing in two places)[37].

Our system can formalize also the statement: initially, there were not two ships, only one; at the end, two ships exist but they are not identical to each other, though each is in some sense "the same" as the original: one same-by-continuity, one same-by-matter. We have two relations to capture those: classical identity (for continuity case) and ≡_π (for matter case). Because we allow multiple modes, we avoid any contradiction from "two ships are the Ship of Theseus" – we clarify: one is the same ship strictly; the other is the same collection of matter but not the same ship strictly. Leibniz's Law, again, is upheld: the museum ship has the property "underwent repairs gradually", the reassembled has "underwent disassembly and storage", so they can't be identical because those properties differ.

This is a nice demonstration of how π-identity can encode *principled* answers to identity puzzles without appealing to brute intuitions. We essentially encode the criteria of identity for artifacts (continuity vs original parts) and show that when they diverge, so do the kinds of identity we talk about. There is no formal contradiction in saying "Ship A = Ship original, Ship B ≡_π Ship original" – since ≡_π is a different predicate, the statements can coexist.

One could choose a different criterion: Some might argue the real identity lies in the planks, not the continuity. Our framework could accommodate that by flipping what counts as non-branching. The point is, we can represent either criterion clearly; but importantly, trying to have both as "=" leads to inconsistency, which we reveal, forcing a choice or a layered answer.

(Optional 7.4 would discuss identity in software or information, but given space, we'll skip a detailed write-up.) It would likely say: A machine learning model's identity might be tracked by provenance of training data and weights. If you clone a model's weights and then fine-tune separately, are the two resulting models identical? Not strictly (they diverge after cloning), but they share provenance. Our logic would treat them like fission: original model's lineage branched into two weight-histories. We'd have no strict identity after, but perhaps a defined metric of lineage identity (like identical up to checkpoint X). This again highlights the use of ≡_π.

8. Theology Without Special Pleading

We turn to theological applications, showing that the same logical system L_π that worked for statues and ships also handles profound mysteries of Christian doctrine in a natural, disciplined way. The key move is to separate **predication (∈)** from identity and to use **provenance labels** to mark the relations of origin that distinguish divine persons. Classical theology often resorts to phrases like "one in essence, three in person" or "without confusion, without change" – we will give those a formal home in PEI.

8.1 Trinity: Consubstantiality Without Collapse

The Doctrine: There is one God (one divine essence or substance), in three Persons: Father, Son, Holy Spirit. The Father is God, the Son is God, the Spirit is God – yet the Father is not the Son, etc. They are distinguished by their relations of origin: the Father is unbegotten, the Son begotten of the Father, the Spirit proceeding from the Father (and, in Western doctrine, from the Son also) – these relations do not partition the substance but are *subsistent relations* within the one Godhead[45].

The Logical Problem: Naively, "Father is God" and "Son is God" look like F = God and S = God, which by transitivity would imply F = S, a heresy (Sabellian modalism). Another attempt: treat "is God" as predication of a nature. Then we have F ∈ Divine, S ∈ Divine, which is fine. But we also need to say "There is exactly one God" which sounds like "there is exactly one divine substance or nature". How to assert oneness of God without identifying the persons? Some use relative identity: "the Father and Son are the same God but different Persons" – but "same God" in relative identity is a non-standard logic many find problematic[6][31]. Our goal is to avoid relative identity in that sense, while preserving the content: they share one nature (same God) but are not identical as persons.

PEI Resolution: We start by representing each Person as a supposit with the Divine nature: `Father ∈ DivineNature`, `Son ∈ DivineNature`, `Spirit ∈ DivineNature`. Now, DivineNature is special: it includes uniqueness (there is only one Godhead, one divine

essence). We can represent the unity of essence in two ways: (i) treat DivineNature as an abstract entity that can be "multiply instanced" by persons but remains numerically one (like one universal wholly present in each – akin to Scholastic formal distinction), or (ii) we could go so far as to say the three persons somehow share the same esse (act-of-being), essentially the same individual Godness. Aquinas for example would say they are numerically one in essence[46]. If we wanted to formalize *that*, we could introduce an entity "God" or "the Divine Essence" as a supposit or something, but better not – we don't want a 4th thing. Instead, we say: whenever P ∈ DivineNature and Q ∈ DivineNature, then P ≡_E Q (some equivalence meaning same existence or same being). Perhaps we can use our ≡_π cunningly: In God, origin (π) distinguishes persons, not substance. So what if we assign: π labels for Father, Son, Spirit as U (Unbegotten), F (from Father), FS (from Father and Son). Then consider the possibility that *esse* is indexed by not person but by nature: a single act-of-being E_G that all three "use". Our system can express that by saying something like: E(Father) = E(Son) = E(Spirit) in the model (one existence). But we won't go too far – suffice: they are consubstantial means they are all ∈ DivineNature and there's exactly one such nature instance. A formula: ∃x∀y (y ∈ DivineNature → y = x). However, that would identify Father and Son by that x, so scratch that approach. Instead: ∃!N (N = DivineNature ∧ Father ∈ N ∧ Son ∈ N ∧ Spirit ∈ N). This asserts one shared nature concept.

Now, we need to avoid F = S. We do that by giving them distinct provenances: define π(F) = "U" (no origin), π(S) = "F" (from Father), π(Sp) = "FS" (from Father and Son as one principle). Then we add an axiom: If two divine persons had the same origin, they would be the same person. Here, none have same origin pair (Father has none, Son has Father, Spirit has Father+Son). Thus they remain distinct. We actually treat these origin tags as we treated branching: the only way to have two distinct persons with one divine nature is to have some "branching" at the eternal level. But that's not branching in time, it's the personal processions. We incorporate it by saying: in the Trinity, the π is not a temporal process but an ontological ordering: Father's π is unique, Son's π indicates sourced from Father, etc. Our logic doesn't natively have a notion of eternal generation, but we can treat it as analogous to an origin relation.

Within L_π, we can simply *assert* non-identity: Father ≠ Son, etc., and those are taken as givens (we are consistent in doing so). We also assert the analogy: Father ∈ Divine, Son ∈ Divine, etc. Now consider any attempt to prove something heretical like Father = Son. Could it be derived? Only if some rule equated them. The likely culprit would be substituting their common "is God". But "is God" we represented as predication, not identity. So Father ∈ Divine and Son ∈ Divine does not allow Father = Son by any rule; that would be like saying "F and S are both in set {God} so F = S", which our logic disallows unless we had an axiom "at most one divine supposit" which we did not assert. We only assert one divine nature, not one supposit. So already, by design, the collapse cannot be derived.

We should ensure not to inadvertently allow "one essence = one person" confusion. For instance, earlier we avoided ∃x∀y (y ∈ Divine → y = x). Instead, we reflect the creed: "They are one God because they are one in essence" (Augustine's phrase: Father and Son "are

one wisdom because one essence"[47]). In our terms: Father and Son share all essential properties (because all those are properties of DivineNature), but they differ in relational properties (Fatherhood, Sonship). And since substitution is restricted, one cannot substitute "Father" for "God" in a statement like "God is one". Perhaps we express "God is one" as: any two divine persons are ≡ with respect to essence or something.

Another important feature: The Trinity's persons are distinguished purely by relations of origin, which are *incommunicable* properties (Father alone is unbegotten, Son alone begotten, etc.). We represent these as predicates like Begotten(x,y) etc., or just in π labels. Perhaps simplest: a predicate FatherOf(x,y) for x fathers y; then we have FatherOf(Father, Son), and either directly something for Spirit. But we need not formalize theology fully; the main point is: identity is downstream of origin, and here origin differences yield personal non-identity despite an essence unity.

We can attempt a **consistency proof sketch**: Assume all doctrinal statements in L_π: - F, S, H are distinct supposits. - F, S, H all ∈ Divine. - Divine has analogical predicates like almighty, eternal, etc. So Almighty(F), etc. - There's exactly one almighty being (but since almighty is said analogically of each person, no conflict). - Relationship: Orig(S) = F, Orig(H) = {F,S}, Orig(F) = none. And maybe a rule: if Orig(x) = Orig(y) = none, then x=y, guaranteeing only one unoriginate (this fits: only Father unoriginate). We have no conflict as only Father meets that, so fine.

From these, we try to derive a contradiction. Where could it come? Possibly someone says: "Father is God; Son is God; so Father = Son." But in our formalization, "is God" is not an identity statement, so the inference is blocked. Or: "There are three who are God; so either three Gods or one of them is not fully God." We answer: there are three who are God by predication, but "three Gods" would mean three distinct divine substances. In L_π one could define "x and y are same God if x ≡_π^E y" meaning share same act-of-being. If we assume all three share one esse (which is a theological assertion of simplicity), then Father ≡_π^E Son ≡_π^E Spirit (some equivalence on being). And maybe we'd say if two persons share the same esse, they constitute one God (like one being). That doesn't identify them, it just says they are not separate beings. This is how we can logically capture "one God in three persons".

So formally: ∀x,y (x ∈ Divine ∧ y ∈ Divine → x ≡_E y). Then add rule: if x ≡_E y, then any "being-level" predicate holds jointly. For example, "there is one omnipotent being" can be true because F, S, H are all ≡ on being and omnipotence is a being-level attribute, so you don't count it thrice. Without diving deeper, it's plausible that no contradiction arises because we keep track of the category of predication (person vs nature).

Thus, we have solved the logical problem of the Trinity by the classic approach of distinguishing *is* of identity from *is* of predication (something the medievals implicitly did) and now formally enforcing it with substitution rules. Our contribution is to incorporate the **origin constraints**: we can formalize "Father and Son are not distinct in any way except their relation of origin"[21][48]. In L_π: all properties except those that involve provenance come out identical for F and S. Indeed, if a property P is about essence (like "almighty,

creator, eternal"), then from F ∈ Divine and S ∈ Divine, P applies to both; and since those are in extensional context, one might even say P(F) and P(S) and thus they are "indiscernible" in essence. Yet since π(F) ≠ π(S), one property differs: "being unbegotten" holds of F not S, and that's a π-indexed relational property, not an essence property. That suffices to keep them distinct persons. So indiscernibility of identicals holds: they are not identical because there is at least one discernible difference (origin). But they are indiscernible with respect to essential attributes, reflecting *consubstantiality*. This matches Aquinas: "Everything is one [between Father and Son] where there is no relation of opposition"[49], i.e., all non-relational attributes are shared.

Thus our system yields exactly that: relative to non-provenance predicates, F and S are indistinguishable; only relational predicates (like has-origin-from-someone) differentiate.

8.2 Christology: One Person, Two Natures

The Doctrine: Jesus Christ is one person (the eternal Word), who has both a divine nature and a human nature, "without confusion, change, division, or separation" (Chalcedon, 451). Because of this, statements like "Jesus is God" and "Jesus is man" are true. Moreover, due to the unity of the person, properties of either nature can be predicated of the person (communication of idioms): e.g., "God [the Son] suffered and died on the cross" is in orthodox theology a true statement – not that the divine nature suffered, but the Person who is God suffered in his human nature. Conversely, "the man Jesus is omnipotent" is also somehow true because the subject "the man Jesus" refers to the Person who, in his divine nature, is omnipotent. However, one must not mix up the natures: one should not say "Jesus as God is passible" or "Jesus as man is eternal" – those cross attributions are false if meant literally, and the natures remain distinct.

The Challenge: Classically, logic would say if Jesus = God (the Son) and Jesus = man (a human being), then God = a man – which seems either heretical or nonsensical (divine nature equated to a creature). The communicatio idiomatum appears to violate predicate logic: "God died" (subject God, predicate died) sounds false because God in himself cannot die; yet Christians assert it true in context ("God the Son died in the flesh"). How to rigorously allow these mixed predications without confusion?

PEI Resolution: We already have the tools. We represent **Jesus as one supposit with two natures**: let J be the supposit designating the Person of the Son Incarnate. Then we assert `J ◇ ⟨DivineNature, HumanNature⟩`. This single statement encodes the hypostatic union (two natures united in one hypostasis/person). From it, we can derive `J ∈ DivineNature` and `J ∈ HumanNature`. Now, consider any predicate: e.g., "eternal" is a property of the divine nature; "born of Mary" a property of human nature. We reflect this by either indexing predicates or by context: We could have a rule that if N is DivineNature, from J ∈ N infer Eternal(J) (given Jesus is eternal in his divine nature), and if N is HumanNature, from J ∈ N infer BornOfMary(J) (he is born in the human nature). So we get Eternal(J) and BornOfMary(J). But now if someone tries to infer something like "Mary is the mother of God" – actually that one is a doctrine ("Theotokos": Mary is Mother of God because the

person she bore is God). In our system: Mary being mother of J (as man) along with J ∈ Divine means Mary is mother of a person who is divine – we could allow a derived predication "MotherOf(Mary, J)" and from J ∈ Divine maybe infer a term "MotherOfGod(Mary)" if carefully defined. The system doesn't directly yield that phrase, but it's expressible: Mary is mother of the one who is God, so in ordinary language we call her Mother of God – nothing logically amiss if parsed correctly.

Now, **substitution discipline** ensures no confusion: We cannot substitute from one nature context to another. For example, from J ∈ DivineNature we get "Immortal(J)" (cannot die) in context of divine nature; from J ∈ HumanNature we get "Mortal(J)" in context of human nature. These are not a contradiction because they are tagged by different modes. If someone dropped the tags and said "J is immortal and J is mortal, contradiction!" that person would be mixing modes and our logic forbids that inference. It's exactly like the earlier Jesus omnipotent/mortal reasoning we did. So **no contradiction** arises.

How do we handle communicatio idiomatum properly? E.g., "God died on the cross" – here subject "God" refers to the Second Person (not the Father or Spirit, typically one clarifies "God the Son"). In our formalism, "God the Son" is J in his divine nature. The statement is effectively predicating a human attribute of a divine subject because of the unity of person. One way: since J denotes the person (who is one), we just say "Died(J)". We know died is a human nature predicate. By aspect-restricted substitution, can we assert Died(J)? We have J ∈ Human, which gives Died(J) (allowed in human channel). Now the question: can we also say J ∈ Divine and still say Died(J)? The rules said do not transfer predicates across nature channels. But the communicatio idiomatum is precisely doing that in language – saying something proper to the human nature ("died") of the person who is divine. Our system must allow that *statement* while understanding it's true because of the human nature. How to do this logically without breaking the rules? Possibly we treat such statements as abbreviated forms: "God (who is J) died" is true because J (in his human nature) died. So as long as the listener knows that, it's fine. Formally, we might handle this by context or by introducing a notion of qua predicate: "Died(J qua human)" but one usually doesn't put that in everyday talk.

We might incorporate a special rule: if x ◇ ⟨N1,N2⟩, one nature is substantially higher (divine) and a predicate applies in the lower (human) channel, one can assert it of the person without qualification. This is a *communicative rule*: **Communicatio Rule:** If x ◇ ⟨N_divine, N_human⟩, and P is a predicate applicable to N_human (e.g. "died"), then we allow a derived statement "P(x)" to be asserted (with the understanding it's true via the human nature). We must be careful not to allow the reverse: a divine property asserted of the human nature part in a way that implies created nature has that property by itself. E.g., "Jesus's human nature is omnipotent" – we should block that because omnipotence is only via divine nature. Indeed our substitution rules already block substituting J into an "is omnipotent" context under human aspect. The communicatio idiomatum works in one direction: from nature to person. The person can be referred by either nature and you can state properties from either side of his being of that one person. But you cannot attribute

properties of one nature *to the other nature itself*. Our system handles that: we can say Person J (who is divine and human) has property from either nature. But we cannot say "DivineNature(J) implies property of HumanNature(J)" or vice versa, since that would require substitution across aspects. So the upshot: "God died" is fine (J has divine nature so calling him God is predication, and died we know via his human nature, we state it of J), "A man is omnipotent" meaning J as man is omnipotent is also fine if we mean J (the person) has omnipotence by virtue of divine nature – though usually one wouldn't phrase it that way.

Thus, the framework exactly captures Chalcedon's formula: the natures are distinct (we never equated them or mixed their properties in one soup), but the person is one (J is one supposit). **No contradiction** (like mortal and immortal in same subject in same sense) occurs. The formal regime of aspect substitution is essentially a bulwark against Nestorian splitting (treating as two persons) on one side and Eutychian confusion (mixing natures) on the other – in logical terms, Nestorianism would be treating J as two separate subjects which we didn't (we have one J), Eutychianism would be allowing cross-substitution as if one nature and we did not. So orthodoxy is maintained naturally by the logic.

We could illustrate with a mini-proof:
Premise: J $\diamond$ ⟨Divine, Human⟩.
From this, `J` $\in$ `Divine` and `J` $\in$ `Human`.
We have axioms: $\forall x$ (x $\in$ Divine $\rightarrow$ Eternal(x)), $\forall x$ (x $\in$ Human $\rightarrow$ Born(x) $\wedge$ Mortal(x)).
Thus Eternal(J) [via Divine], and Mortal(J) [via Human]. According to our logic, Eternal(J) and Mortal(J) are not contradictory because their proofs carry aspect tags. In fact, if someone tries to derive a contradiction: assume Eternal(J) and ¬Eternal(J). The latter would follow from Mortal(J) only if there's some axiom linking mortal $\rightarrow$ ¬eternal. But mortal in context means "will die or can die", eternal means "exists outside time or forever". Strictly, a mortal being is not eternal. But here J is mortal *in his human nature*, eternal in his divine. We cannot combine these without error. If our logic had a global axiom "$\forall x$ (Mortal(x) $\rightarrow$ ¬Eternal(x))", that would cause trouble by ignorantly mixing. We wouldn't include that globally; we might index it: $\forall x$ (x $\in$ Human $\rightarrow$ ¬Eternal(x)), which is true for any purely human supposit. But J is not purely human; J $\in$ Human, yet J is also eternal by another nature. Our logic doesn't automatically apply that rule to J, because the rule's antecedent is that x is human *and nothing else*. Actually, x $\in$ Human and not also Divine might be needed. In practice, we avoid such general axioms or we treat them carefully. If done right, no inconsistency arises.

Thus, the calculus L_π shows that far from being an incoherent mystery, the Chalcedonian two-natures doctrine is logically consistent and well-handled by a typed modal logic of identity. It needed no exotic "relative identity" trick – just careful typing of predication.

8.3 Grace and Adoption: New Provenance in Creatures

The Doctrine: In Christian theology, a human person who receives sanctifying grace becomes an adopted son/daughter of God, a partaker in the divine life (2 Peter 1:4). This is

not a substantial change (the person remains human, not becoming divine by nature) but a real relational and qualitative change: a new created disposition (habitual grace) and a new relation of filiation to God through Christ. Grace is said to make us "children in the Son" – our sonship is secondary, through the true Son.

PEI Application: This scenario beautifully fits our provenance notion. We can say: originally a person P has provenance π = "born of human parents X and Y at time t". After baptism (say), P acquires a new, supernatural provenance tag as well: being born of God in a spiritual sense. We might formalize that P now has a π that references Christ's lineage: Christian theology often says we are sons of God by adoption, meaning we share in the Son's sonship (through the Holy Spirit's action). We could represent that $P \equiv_\pi{}^{spiritual}$ Christ in some analogical way: not literally same origin, but we tie P's new life to Christ's merits (this might be too theological for logic, but one can imagine an abstract model where all the saved have a pointer in their π "connected to Christ").

The main point: **no change of identity** of P is asserted. P remains the same supposit (no branching, just an added aspect or relation). P certainly does not become identical to Christ or to God. However, P gets a new status: P ∈ Graceful or P ∈ Adopted as a nature-aspect? "Adopted child of God" is not a new substantial nature (not like becoming an angel or something), it's more of an accidental but supernatural form. We could treat it as a secondary nature in a broad sense: perhaps `P ◇ ⟨HumanNature, AdoptedChildOfGod⟩`. The second is not a substance but a state. Alternatively, treat adoption as just a predicate.

Anyway, the logic helps by clearly demarcating: our deification (becoming "like God") is via a relation and gift, not by essence. We will not have confusion like Mormon theology (where they might say men can become Gods literally – that would be an identity or nature change, which in orthodox view doesn't happen). In L_π terms: no human gets ∈ `DivineNature` by grace; they remain ∈ `HumanNature` but gain some new Int or π connection maybe. This ensures that while we speak analogically of them as "children of God," we don't violate the metaphysical gap that Creator and creature remain distinct in essence.

So in short: Grace introduces a new mode of identity or union – not identity of being (we don't become the same being as God or Christ), but a provenance link (like branch grafting: Christ as vine, we as branches, to use the biblical metaphor). Provenance in Christ means our destiny now follows a new line (eternal life) as opposed to mere natural lineage (which leads to death). In logic: perhaps before, human H had π = natural lineage. After grace, H's π is augmented or partially merged with Christ's π regarding spiritual life. But not fully sure how to formalize, we can say: an extrinsic denomination: God calls H a son now. That is a new relation R(H, God) true.

Our system doesn't directly add much here except to highlight we are adding a relation, not altering identity or nature. The discipline of substitution prevents one from saying "Since John is a son of God and Christ is the Son of God, John = Christ (or = a son by nature)" – nonsense avoided by understanding "son" in different senses (analogous vs univocal). We could formalize analogy by context but that's beyond scope.

8.4 Sacraments and External Unity of Action

The Doctrine: In Catholic theology, sacraments (like the Eucharist) involve changes that are mysterious: e.g., bread and wine become the Body and Blood of Christ (transubstantiation), but without any observable change in accidents. Also, when God acts in the world (as in sacraments or any grace), the three Persons act together (indivisa ad extra[50]) – there is no separate act of the Father vs Son vs Spirit toward creatures, they all concur as one cause. However, in the economy, certain actions are appropriated to one Person (we say the Son became incarnate, not the Father), yet all divine actions outwardly remain one.

PEI Resolution: We can sketch two points:

1. **Transubstantiation:** Our logic can handle the Eucharist by using modes: The Church says the *substance* (essence) of bread is changed into the substance of Christ's Body, while the *accidents* (appearances of bread) remain. In L_π, we could have at consecration: one supposit x that was Bread now becomes x ∈ `BodyOfChrist` and ceases x ∈ `Bread`. Not a usual event, but integrity of bread fails (since substance changed), a new integrity (body of Christ present) obtains. Is x the same x? Classical doctrine says no substance of bread remains; so the bread's identity is gone, replaced by Christ's presence. We represent: original bread B with π_bread; after, we have entity C (Christ's body piece) with π perhaps pointing to Christ's historical body origin. It's almost a miraculous *fusion*: Christ's substance becomes present in place of bread's. Our system forbids saying bread = Christ, because clearly different provenance (bread from wheat, Christ from Mary). But at the moment of consecration, bread's substance is annihilated (in a way) and Christ's substance becomes present sacramentally (without a new local motion from elsewhere). We'd treat that as a supernatural reassignment of what supposit occupies that location. We might not formalize this fully, but conceptually: If one tried to say the bread's accidents are predicated of Christ's body (like Christ's body appears as bread), our aspect system allows one supposit (Christ's presence) to have the accidents of bread by a special divine appointment. This is unique, but it uses again the idea of multi-aspect: Christ's Body present under aspects of bread. So y ◇ ⟨BodyOfChrist, BreadAccidents⟩. And no contradiction because those aspects are quite different categories (substance vs accidents). This prevents naive contradiction like "this is essentially Christ and essentially bread" – no, it's essentially Christ, only accidentally bread in appearance. Formal logic can support that by distinguishing essential vs accidental predication and disallowing mixing them as if one.

2. **Divine Action (indivisa ad extra):** In our framework, one could formalize that any action on creatures has one and the same origin from the Trinity. We could say: any effect in creation E has a provenance label that includes all three Persons as co-cause. Yet we appropriate within that common cause a distinct relational aspect: e.g., the Father sends the Son, the Son is incarnate, the Spirit sanctifies, etc., but

each of those events still has all Persons behind it in unity. How to notate: We might treat the Trinity as one agent externally (one principal cause, like one divine will). So an effect can be labeled π_effect = "from God". Meanwhile, within the internal discourse, we can say the Son did this or that, but that's an aspect of the one divine action (like channel Father, Son, Spirit for the one external action). Since our logic easily allows multi-aspect for one supposit, we can imagine the "agent" in a miracle is G ◇ ⟨Father, Son, Spirit⟩ in terms of roles. But that might be pushing it.

At least, we can ensure not to derive false things: e.g., God's actions are one doesn't mean Father = Son (we already handled that). Or saying "the Father alone created" (heresy) is avoided because we'd formalize creation as an act of the divine essence, which all Persons share, so our logic would reflect all Persons ∈ DivineNature yield their involvement in creating.

In summary, the theology cases show that by combining classical identity for the core substance with new modal aspects for relations and predications, we can formalize mysteries in a way that is logically consistent and actually aligns with the historical orthodox explanations (which often relied on analogical predication and distinctions like person/nature, substance/accident – precisely what we've built into L_π). The payoff is *one unified logical framework* where secular and sacred puzzles alike are solved by the same principles: identity is nuanced by types and provenance in both realms.

9. Historical Positioning and Comparison

The π-identity calculus and PEI paradigm consciously honor and extend a classical worldview: **metaphysical realism** (there's a real structure to being that our logic should track), and the idea that *being is first, identity second*. Historically, thinkers from Aristotle through Aquinas did not fetishize identity as a primitive; they cared about what makes a thing one (form, act) and what differentiates it (matter, relations)[23][45]. Our system essentially resurrects that approach in formal guise.

We can frame this as a **conservative paradigm shift**. It's conservative because, in the core extensional fragment, we have done nothing to contradict classical logic – any equation that was provable before is still provable (Conservativity claim proven in Section 6). And any classical valid argument not touching the new modalities remains valid. Yet it's paradigm-shifting in that we **reject the modern default** of treating identity as a one-size-fits-all relation. Instead, we offer a refined toolkit: absolute identity exists but is only applicable within well-defined boundaries; outside those, we use new relations (constitution, provenance equivalence) or layered typing.

Comparison with Relative Identity Theories: Geach's relative identity (RI) suggested that statements of identity are incomplete without a sortal – effectively that there is no single absolute identity at all[6][7]. Our approach differs crucially: we **keep absolute identity** ("=") as meaningful and in fact as the final word when applicable (the highest bar of sameness). We do, however, incorporate sortal information into our logic by preventing cross-sortal confusions. In a sense, we realize the spirit of RI (that "same F" vs "same G"

can both apply to two things differently) but without abandoning absolute identity – we just say in those cases of puzzle, absolute identity doesn't hold at all; something weaker does. Geach would say perhaps the Persons of the Trinity are the same God but not the same Person – sounding like relative identity. We say: Father and Son are two distinct supposits, each ∈ DivineNature (so each "is God" in predicative sense), but you cannot infer one supposit from the other due to the rules. This yields essentially the same inferences as RI in practice, but we maintain a single equality symbol and logical orthodoxy in core.

In sum, we **avoid the pitfalls of RI**: RI frameworks often face problems of coherence (how can identity be equivalence within each sortal but no overall equivalence?). Our approach doesn't treat "≡_π" as identity, just as a weaker relation. So we don't say "same God" in a logical primitive sense; we say "two persons share the one divine nature" which is a predicate fact, not an identity relation. This is less philosophically controversial because it's like sharing a property rather than bending identity.

Comparison with Sortal Essentialism (Wiggins, Lowe): Philosophers like Wiggins (1967, 1980) insisted that every identity statement is under a sortal concept, and that objects have sortal essences determining their persistence[51][26]. Our Integrity conditions Int_F directly model that idea: to know if a = b later, you need to know under what sortal you track it. We go beyond Wiggins by formalizing not just the criterion but also what happens when an object falls under multiple sortals or changes sortal (Wiggins generally would say true sortals can't change for an object – if it seems to, it's a different object of another sort; our system can emulate that by ending identity at change). We also integrate sortals with modalities like provenance. Sortal logic as usually presented (e.g., in E.J. Lowe's four-category ontology or others) doesn't have an explicit mechanism for things like fission or co-location beyond saying it's not possible or leaving it at intuition. We provide a concrete rule ("non-branching or no identity") which enriches the sortal framework with a temporal/causal dimension. So one could say we stand on the shoulders of sortal essentialists but extend their work to handle previously troublesome cases (like the statue and lump – which Wiggins handled by denying identity and using constitution language we mirrored[27], or the Trinity which sortal logic alone wouldn't solve since persons share a sortal "God" in common yet are distinct – our addition of π-relations covers that missing piece).

Comparison with Truthmaker/Hyperintensional Theories: Our system is hyperintensional in that it can distinguish different ways in which a statement can be true (because substitution can fail even when classical truth conditions are same). For example, two formulas "a is necessarily F" and "a is contingently F" might be extensionally equivalent for a certain a, but our logic would keep track of the mode. In truthmaker semantics (Fine, etc.), one would say we need to consider the exact truthmaking state. Similarly, our provenance is a bit like a "truthmaker" for identity statements – it's the condition that must obtain to make x=y true. When it fails, only a weaker statement holds. We can cite Fine's remark that a hyperintensional concept allows distinctions beyond logical equivalence[52]. The Trinity example shows hyperintensionality: "Father is God" and "Father = God" might be extensionally equivalent in models where only one being is

God, but we differentiate them intensively. Our approach is in line with current hyperintensional logics but specifically tailored to metaphysical identity conditions, rather than, say, propositions or content.

Constitution Views in Metaphysics: There is a literature (e.g., Lynne Baker's Constitution theory) claiming constitution is a relationship distinct from identity that accounts for cases like statue and lump. We are very much in agreement with those views[1]; what we do is provide a formal rule-based account of constitution (our ◇ and Int rules). Baker would say the statue and clay are two distinct things in a "unity relation". We formalize unity via shared matter provenance and aspect substitution constraints.

Truthmaker and Essence (Fine): Fine's work on essence (e.g., distinguishing between essence and modality) resonates with our Esse vs Integrity separation. We explicitly have E(x) as fundamental existence principle, which was a key in Scholastic thought and revived by Fine's focus on essence needing its own tools besides just modal logic[53]. We incorporate essence indirectly by Int_F (which is essentially capturing what is essential to being F).

In summary, the PEI/π-identity framework is **classically rooted** but novel in execution. It doesn't ask metaphysicians to throw out standard identity or classical logic (in fact, it requires them to keep classical logic at the core), but it *augments* it in a way that addresses both modern puzzles and respects ancient wisdom. As a result, it's a **unified discipline**: where relative identity was a radical redefinition, and sortal logic was a partial measure, and type theories stayed mostly out of metaphysics, we combine elements of all into one coherent system.

10. Objections and Replies

We anticipate several objections:

Objection 1: "Provenance-identity (≡_π) is just relative identity in disguise." One might say we effectively allow things to be "same in one respect, different in another," which was exactly Geach's idea of relative identity – so aren't we falling back into that, with more notation?

Reply: Not exactly. In relative identity theory, one rejects absolute identity entirely, permitting statements like "x and y are the same F but not the same G" without any underlying notion of x and y simply being the same thing[6]. In π-calculus, we **never say "same simpliciter" unless we mean it absolutely**. If ≡_π is true, it does **not** mean x and y are partially identical entities; it means they are distinct entities with a particular common origin relation. We keep a clear distinction between "=" (which still obeys all classical laws) and "≡_π" (which doesn't). There's no confusion of identity criteria because "=" remains one and transitive; ≡_π is another relation entirely. We thus avoid the logical paradoxes that plague relative identity theories (like the problem of "relative Leibniz's Law" etc., which Geach's theory struggles with[30][54]). Our ≡_π is more akin to what some call an equivalence relation or a similarity relation, not identity. When we say "the Son and

Father are the same God" in plain language, our formal interpretation is not an identity predicate at all but "Son ∈ Divine and Father ∈ Divine" plus maybe "DivineNature is unique" – a predicative fact. So we carefully sidestep the entire commitment relative identity makes (that identity is many-headed). Instead, we say identity is one-headed, but sometimes you're dealing with something that's not identity at all, just a sibling relation often conflated with identity.

Objection 2: "Why introduce new modes? Couldn't one use existing multi-sorted logic or type theory to do the same?" For instance, maybe we could treat human and divine nature as types in a typed lambda calculus or something, or use second-order logic to enforce sortal constraints. Are our modes really adding anything or just restating sortal logic in a complex way?

Reply: It's true we draw inspiration from type theory – one can see ∈ N as a type assertion. But normal multi-sorted logic usually doesn't allow one constant to have two sorts at once; we specifically needed that for Christ (one entity of two natures) and similar cases. Dependent type theory or type hierarchies might allow that (like subtyping). However, our system is more coarse-grained and metaphysically informed: it distinguishes provenance line (something not typically in type systems) and restricts substitution differently than a purely type-based approach would. Typical sorted logic could prevent saying "this number = this person" because of sort clash – that's fine but trivial. Our concern was preventing more subtle mistakes like substituting a term in an intensional context across possible worlds or across "qua" contexts. That goes beyond standard type theory. Also, we present the system in a user-friendly logical calculus rather than expecting metaphysicians to adopt full dependent type theory notation. In a sense, we're packaging some of the benefits of advanced type theories (like keeping track of categories and preventing category errors) in a first-order logic shell that is more familiar. So the modes are not redundant; they are the front-end for those who might not want to deep-dive into type theory but need its discipline.

Objection 3: "Does branching identity introduce vagueness or arbitrariness? If identity can cease when 'branching' occurs, isn't it sometimes vague whether something is one object or already splitting into two? Could it lead to indeterminate identity?" For example, in fission, at what exact point do we say identity ended and two began – could that be fuzzy?

Reply: In many puzzles, yes, if the scenario has a gradual divergence, there might be vagueness. But that vagueness is in the phenomenon, not in our logic. Our logic will say: as long as Int_F holds and there's no branch, identity holds. Once clearly branch, identity fails. If a situation is genuinely indeterminate if/when a branch occurred (like psychological divergence might be gradual), then indeed whether identity ended might be indeterminate – but that's appropriate because the world's description is indeterminate. Our logic can handle that by not giving a definite verdict in that case (perhaps both identity and non-identity might be non-theorems unless further stipulation). However, usually thought experiments specify a clear fork. Also, our formalism can be sharpened or relaxed: we

might define branching in a crisp way – e.g., "two distinct continuers at time t1 = branch at t1." That's quite precise. If one wants to accommodate vague branching (like in quantum identity issues), one could integrate a supervaluation or fuzzy logic layer atop our system, but that's beyond current scope. In summary, we don't think we've introduced new vagueness; we've just acknowledged that certain identity attributions are non-binary in extreme cases – which most metaphysicians already accept (think of debates on the exact point a growing heap becomes "not the same heap").

Objection 4: "Aren't you smuggling theology into secular metaphysics? Why should a secular metaphysician accept these new primitives (π, etc.) just because they make Trinitarian doctrine work?" The worry is our motivations might be seen as theological and thus not compelling for those uninterested in theology.

Reply: We certainly were guided by theological use-cases, but the framework stands on its own merits for secular cases. In fact, we first demonstrated it on secular puzzles (Section 7). The ability to handle Trinity or Incarnation is a bonus that shows the framework's power and coherence across domains, but one need not be religious to appreciate the solution to the Ship of Theseus or fission paradox. Provenance and integrity are neutral, natural concepts: origin and persistence conditions. These were lurking implicitly in many secular debates (e.g., "only one of the fission successors is *the* original because maybe original's cause line can't bifurcate" – a secular reasoning we've just formalized). So nothing essentially theological is assumed in the core logic; theology only inspired the multi-nature and analogical predication aspects, but those have analogues in secular ideas of multiple categorization (e.g., an object can be a member of multiple sets or categories without identity collapse). We did ensure that our formal apparatus aligns with Catholic orthodoxy, but that only shows it's robust—if it can respect such strict constraints, it's likely consistent and careful enough for any domain. So secular metaphysicians can use π-calculus just to solve material constitution, personal identity, modality issues, and ignore the theology sections if they like. In doing so, they lose nothing and perhaps gain a unified view that outperforms patchwork solutions (like having one logic for coincidence, another ad hoc approach for fission, etc.—here it's one system for all).

In short, while theology was a testbed and beneficiary of our calculus, the calculus is not inherently theological. It's metaphysical and logical. Theology just happens to be one field where intensional identity issues are stark, so it provided good validation.

11. Extensions and Open Problems

The π-identity calculus opens several avenues:

- **Categorical Semantics:** It would be fruitful to recast L_π in category-theoretic terms. Possibly, we could interpret provenance as morphisms or path-connectedness in a category of objects. A categorical semantics might treat the worldlines (provenance) as morphisms from a timeline category to an object category. The non-branching condition might relate to monomorphisms or

something similar. This could clarify how identity behaves in abstract settings (like in categorical set theory or in homotopy type theory, where identity is a path).

- **Connections to Dependent Type Theory:** Indeed, some aspects of our system resemble dependent types or multi-modal type theory. We could try encoding PEI in a proof assistant (Coq, Agda) by using types for natures and perhaps an indexed monad for provenance. Dependent types might naturally enforce some substitution discipline. Investigating this would make the framework more rigorous and machine-checkable.

- **Temporal and Modal Operators:** We largely spoke about time informally. One extension is to integrate explicit temporal logic (LTL or CTL style) so we can quantify over times and say things like "for all t1,t2: if Int_F(x,t1..t2) and no branching, then x(t1)=x(t2)" formally inside the system. Similarly, adding a modal operator for metaphysical necessity might interplay with π (for instance, origin essentialism claims could be encoded: necessarily, if $x \neq y$ then $\pi(x) \neq \pi(y)$ perhaps, meaning you can't have two distinct with same origin in another possible world – something like that). These additions could further handle puzzles of modality (Kripke's concerns about transworld identity and origin essentialism[20]). Our provenance actually aligns with origin essentialism: an object's origin π is part of its identity, so in any possible world, if that origin condition is changed, it's not the same object. We could formalize such a principle in a modal extension of L_π.

- **Computational Implementation:** It would be interesting to implement a small theorem prover or at least model checker for L_π. Perhaps using Prolog for substitution rules or customizing an existing logical framework. This could allow testing many scenarios automatically and ensuring no contradictions slip in.

- **Empirical Provenance in Law or Science:** Our notion of provenance has a whiff of legal or scientific tracking: e.g., in law, establishing identity of an artwork or person often involves tracing provenance (history, custody chain, DNA lineage). One could apply π-calculus to, say, identity disputes in legal contexts (is this the same ship or car after major repairs? The law has rules – often leaning on continuity or VIN numbers which are like stable provenance tags). Or in science, thinking about identity of a species (common descent as provenance vs morphological identity). These real-world scenarios might provide data to refine what counts as branching or integrity threshold in practice. The framework might be extended with probabilistic elements if needed (for forensic identity: e.g., a clone has same DNA provenance but is a distinct individual – we would say indeed $\equiv$_π perhaps if you count DNA as provenance, but not =; useful distinction).

In sum, the current system could be expanded to a richer logic handling not just a static scenario but dynamic and modal ones, and perhaps interfaced with the rigorous languages of category theory or formal verification for stronger guarantees.

12. Conclusion

We set out to show that **identity, often thought a rock-bottom concept, is better understood as a *derivative verdict* grounded in deeper principles**. Through the Provenance–Esse–Integrity paradigm and the π-identity calculus, we demonstrated a unified approach to many puzzles: from the classic conundrums of statues and ships to the profound mysteries of Trinity and Incarnation. The framework conserves all the successes of classical logic and metaphysics – it never asks us to abandon clear truths like "identicals are indiscernible" in their proper domain[23]. Rather, it enriches our expressive power so that we can say, with formal precision, **why** some things that appear identical are not, and how some things can be "one" in one way and "many" in another without contradiction[46][55].

Key gains include: a clear criterion (non-branching origin under sortal integrity) for diachronic identity that resolves cases like fission without ad hoc stipulations; a logical distinction between predication and identity that dissolves theological paradoxes by respecting the classical person/nature and substance/accident distinctions in formal syntax; and an overall toolkit for "hyperintensional control" – we can distinguish not just necessary vs contingent identity, but the mode in which something is the same or different, a nuance often glossed over.

By addressing both secular and theological cases, we've shown the paradigm's **explanatory surplus**. For secular philosophy, puzzles that needed separate fixes (temporal parts for change, counterpart theory for modality, relative identity for Trinity, etc.) now can be seen as facets of one systematic upgrade. For theology, notoriously seen as resistant to formalization, we provided a disciplined logical language that upholds orthodox positions (no collapse of Trinity or confusion of natures) in a way checkable by proofs, not just analogies.

We close with three *predictions* or at least forward-looking theses that could be tested as research moves forward:

1. **Metaphysical Unification:** The next decades will see a closer synergy between solutions in analytic metaphysics and the structures used in philosophical theology. Our work predicts that ideas like provenance or "intensional identity" will find applications in mainstream metaphysics (e.g., debates on personal identity or quantum particles) once their utility is seen here. If our calculus is right, it should eventually illuminate, say, the debate between endurantism and perdurantism (maybe suggesting a middle ground: objects endure as long as π unbranching, but perdure into branches as separate).

2. **Theological Clarification:** Adopting a formal framework like this will likely clarify and perhaps resolve some intra-theological disputes. For instance, the filioque controversy (Spirit proceeding from Father *and Son*) could be articulated as Spirit's provenance tag being FS vs just F – our system can express "Father and Son are one principle of Spirit" by saying they share giving the origin (one might formalize that as

the Spirit's π has contributions from both). If this framework gains traction, it might show that some debates are purely verbal once formalized (this is a bold hope, but a testable one: try modeling Eastern vs Western Trinity theology in L_π and see if any real logical difference emerges or if it's perspective).

3. **Interdisciplinary Reach:** Identity notions in computer science (object identity in programming, database keys, etc.), in law, and in cognitive science (how we track identity of objects over time psychologically) might all benefit from π-like systems. We predict that a variant of our provenance tracking will appear in those fields. For example, distributed ledger (blockchain) tech already uses something akin to provenance labels to ensure the "same" coin isn't double-spent – essentially preventing branching. If our theory is deeply true, it should resonate outside pure philosophy wherever identity is at issue.

In offering this π-identity calculus, we aim to initiate a disciplined rethinking of identity's role. By treating identity as a *verdict* given certain conditions – conditions which themselves speak to existence, origin, and form – we align more closely with reality's structure as described by both ancient metaphysics and contemporary needs. This is not a denial of the importance of identity, but a repositioning: identity is important *because* of what underlies it (provenance, esse, integrity). And when those underpinnings are accounted for, we can navigate previously dark corners of philosophy with new light and perhaps even find, as we did with theology, that what seemed paradoxical can become transparent under a more fine-grained lens.

In closing, identity was never the ultimate explanatory principle in the classical tradition; it was unity grounded in being. We have formalized that intuition. The "paradigm break" is thus in restoring an older wisdom with modern logical rigor: **Non multa sed unum – not many, but one (in its proper order).**

---

[1] [24] [25] [26] [27] [33] [36] [37] [38] [39] [43] [44] Material Constitution (Stanford Encyclopedia of Philosophy/Summer 2012 Edition)

https://plato.stanford.edu/archives/sum2012/entries/material-constitution/

[2] [3] [5] [9] [10] [23] [32] Identity Over Time (Stanford Encyclopedia of Philosophy)

https://plato.stanford.edu/entries/identity-time/

[4] [6] [7] [28] [29] [30] [31] [40] [54] Relative Identity (Stanford Encyclopedia of Philosophy)

https://plato.stanford.edu/entries/identity-relative/

[8] Aristotle's Metaphysics (Stanford Encyclopedia of Philosophy)

https://plato.stanford.edu/entries/aristotle-metaphysics/

[11] [12] [13] [16] [20] Aquinas: Metaphysics | Internet Encyclopedia of Philosophy

https://iep.utm.edu/thomas-aquinas-metaphysics/

[14] [15] [34] [35] Distinction between Existence and Essence in Avicenna's Ontology and Its Influence on Christian Philosophical Theology with a Focus on Aquinas' Views

https://www.wisdomlib.org/journals/16505-distinction-between-existence-essence-avicenna-s

[17] [47] SUMMA THEOLOGIAE: The persons in relation to the essence (Prima Pars, Q. 39)

https://www.newadvent.org/summa/1039.htm

[18] Negative Attributes | Yeshivat Har Etzion - תורת הר עציון

https://www.etzion.org.il/en/philosophy/great-thinkers/rambam/negative-attributes

[19] Leithart - Page 1189 of 1331 - Patheos

https://www.patheos.com/blogs/leithart/page/1189/

[21] [22] [45] [46] [48] [55] Perichoresis in Aquinas: Fruit, not Foundation | Reformedish

https://derekzrishmawy.com/2017/02/14/perichoresis-in-aquinas-fruit-not-foundation/

[41] [PDF] Time, Fission, and Personal Identity

https://www.argumenta.org/wp-content/uploads/2019/05/Argumenta-John-Perry-Time-Fission-and-Personal-Identity.pdf

[42] [PDF] A Survival Guide to Fission

https://www.uvm.edu/~mmmoyer/papers/Paradox%20of%20Fission.pdf

[49] Whether the Father and the Son are one principle of the Holy Ghost?

https://www.ccel.org/ccel/aquinas/summa.FP_Q36_A4.html

[50] [PDF] Opera trinitatis ad extra indivisa sunt - OPUS

https://opus.bibliothek.uni-augsburg.de/opus4/files/50755/50755.pdf

[51] Sameness and Substance Renewed | Reviews

https://ndpr.nd.edu/reviews/sameness-and-substance-renewed/

[52] Case Study: Hyperintensionalism - Oxford Academic

https://academic.oup.com/book/58055/chapter/477511417

[53] [PDF] 1 Are all primitives created equal? - Durham Research Online

https://durham-repository.worktribe.com/OutputFile/1295814