Particle Species as Representation Classes of Persistent Fold Defects

A Conceptual Reduction of Particle Identity — Scope, Status, and Relation to Existing Accounts

Keith Taylor

VERSF Theoretical Physics Programme

General Reader Summary

The question this paper asks

Physics usually starts with particles. Electrons, quarks, and the rest are treated as the basic ingredients of the world, and the theory's job is to describe how they behave. This programme cannot start that way, because its basic ingredients are not particles at all — they are the Fold *(the smallest possible distinction) and* commitment *(the event that turns a distinction into a permanent record). Particles are nowhere on that starting list. So the question this paper is forced to ask is: if the world is built from committed records, what* is *a particle in the first place?*

The answer proposed is that a particle is not a basic ingredient but a **kind** — a *representation class*. Earlier work in the programme showed that stable bits of matter are knots of persistent structure (called Persistent Fold Defects), and that each such knot is fully described by a short list of properties — its "invariant tuple": things like charge, generation, handedness, and whether it is confined. This paper's central claim is that *a particle species just is one of those property-lists*. An electron is not a tiny primitive object; it is the class of all knots that share the electron's list of properties. Two electrons are "identical" for the simple reason that they have exactly the same list, with nothing left over to tell them apart.

The paper's main argument is that this is not just *a* way to think about particles but, given the rest of the programme, the *only* available one — and the argument has a clean shape worth stating. Take anything you might propose as the "real" carrier of a particle's identity: some hidden substance, a bare "thisness," a secret label. Either it makes some difference you could in principle measure, or it does not. If it makes no measurable difference, it makes no *physical* difference at all, and can be set aside. If it does make a measurable difference, then it simply *is* one more measurable property — another entry on the list — so it has not given you an alternative to the property-list, it has joined it. Either way there is no rival: the property-list is where identity lives. Notice this argument does not need to have *guessed every possible* alternative in advance — any candidate at all, named or not yet imagined, falls on one side or the other of "measurable or not." The one thing it leans on is the assumption that the list of measurable properties is *complete* (that nothing physical hides beneath what can be measured), and the paper marks that assumption plainly, since it is the place the argument could fail.

A second, quieter point may be the most persuasive: physicists *already* identify particles this way. To establish that two things are both electrons, you check that they agree on mass, charge, spin, and couplings — and if they all agree, that settles it; nobody asks "but are they *really* the same?" beyond the measured quantities. That working practice already treats a species as "whatever shares this list of measured numbers," which is exactly the identification the paper proposes. So the paper is not inventing a strange new criterion of particle identity; it is making explicit the one the field has been using all along, and showing it follows from the substrate. As a bonus, the same idea handles unstable particles cleanly — a muon stays a muon for as long as its property-list is unchanged, and "decay" is just the measurable moment the list changes into a different one — so stability becomes a sliding scale, from the electron (whose list never seems to change) to fleeting resonances (whose list changes almost at once), rather than two separate categories.

What this paper is, and what it is honest about not being. This is a paper of *conceptual reduction*, not prediction. It does not calculate the mass of any particle, or how many generations there should be, or any other number an experiment could check. What it does is remove particles from the list of things a theory has to *assume*: it shows that "particle," "species," and "identity" can all be rebuilt from structure already in the programme, without adding anything new. That is a genuine and worthwhile kind of result — comparable in form to realising that "temperature" is just the average motion of molecules, or that "a gene" is a stretch of DNA — but the paper is careful to flag the difference: those famous reductions went on to *predict measurable things that could have been wrong*, and this one has not yet done that. The paper says so plainly, locates itself against the existing philosophical literature on what makes quantum particles identical, and points to the one place the wider programme could eventually make a real, falsifiable prediction — why there are exactly three generations — as the natural next target. It even sketches a candidate answer (three "generation sectors" of cost 1, 2, and 4 exactly fill a closure budget of 7, leaving no room for a fourth of cost 8), but is careful to mark this a *candidate*, because each of its ingredients still needs work: the budget of 7 comes from elsewhere in the programme (a six-sided "closure cell," six edges plus one centre) and *appears* unrelated to counting generations — but the paper flags that this independence has to be *shown*, not just stated, since the whole argument hinges on the two sevens being a real coincidence rather than one built in; and the other ingredient — that the costs double, 1, 2, 4, 8, for any number of generations — is not yet independently established either. So the paper turns "why three?" into a few sharp, checkable questions rather than answering it outright, and is candid that none of them is yet closed.

Abstract

Previous papers in the VERSF programme established that stable matter structures correspond to Persistent Fold Defects (PFDs), that admissible PFDs organise into representation classes characterised by invariant tuples, and that Standard Model gauge structure, flavour structure, and mass-generation mechanisms can be expressed in terms of those classes.

A remaining question is ontological rather than dynamical:

What makes a representation class appear experimentally as a particle species?

This paper addresses that question. It is, by genre, a paper of *conceptual analysis* — an ontological reduction — and not a predictive physics paper: it derives no spectrum, mass, coupling, lifetime, mixing angle, generation count, or exchange statistic, and contains no quantity an experiment could confirm or refute. What it offers is an *identification* together with an argument that the identification is coherent, ontology-subtractive, and consistent with how physics already treats identical particles. That status is stated up front (§1) and isolated again as an explicit limitation (§13A), because the difference between a conceptual reduction and a physical theory is exactly the difference between what this paper does and does not do.

The central result is an identification: stable particle species are representation classes of Persistent Fold Defects. It is stated with its two directions kept apart — that representation classes are *sufficient* carriers of particle identity is argued here (modulo stated conditions: fold-exhaustion, catalogue-completeness, and the restricted Refinement–Transport Bridge inherited from the companion Object Criterion paper); that they are *necessary* is established by a **collapse dilemma** (§11.2A): any carrier of particle identity whatsoever — named or not — either fails to be observable (and so makes no physical difference) or, being observable, collapses by catalogue-completeness into an invariant, hence into the representation class. There is no third branch, so the necessity argument does not depend on having enumerated the alternatives; it rests on catalogue-completeness alone. This is sharper than an elimination of named candidates: it shows even a *successful* carrier becomes a class.

The paper's most defensible claim is stated as a result belonging to physics generally, not to VERSF. The **Operational Species Principle** (§11.4): *given the way particle physics already identifies species* — by measurable properties, admitting no surplus criterion — a species *is* an equivalence class of invariant observables. This owes nothing to the fold substrate; VERSF's distinctive and separable step is only the further identification of that class with a Persistent Fold Defect class. A reader who rejects the fold ontology still owes the general principle, which is forced by existing practice (§6A.1) rather than by the programme. The identification is correspondingly *operational*: every experimental species-identification — measuring charge, spin, generation, chirality, confinement — is already a measurement of the invariant tuple.

The paper's genuine ontological contribution is that the reduction from particle identity down to persistent closure structure is strictly *ontology-subtractive*, removing a primitive at every rung and adding none (§9B), so that the heterogeneous list of particle properties becomes the

structured content of one closure invariant (§9A). A secondary, explicitly *conceptual* consequence — that a species persists over exactly the interval it remains in its class (§4.3) — is stated as an unpacking of the identification, not a theorem; an earlier draft overstated it as a proven biconditional and that overstatement is withdrawn. A species and an object are distinguished throughout (§4.4): objecthood asks *what persistent structure is this individual?*, species asks *which class does it belong to?*

The paper positions the reduction against the existing literature on quantum identity — ontic structural realism, French & Krause's non-individuals and quasi-set theory, Saunders' weak discernibility, and the Hardy / Chiribella–D'Ariano–Perinotti reconstruction programmes (§6A) — claiming unification and substrate-concreteness rather than priority, and flagging where the relationship is unresolved. The closing sections position the result as a *foundation*: the open Standard Model problems of exchange statistics, generation count, and flavour mixing are reframed — as *consequences*, not derivations — into well-posed questions about class structure (§9C). For the generation count the paper goes one step further, offering an explicitly-marked **candidate derivation** (§9.1): three dyadic closure-depth sectors of loads 1, 2, 4 exactly saturate a generation-blind closure register of capacity K = 7, while a fourth (load 8) would overflow it. The candidate's budget premise (K = 7, inherited from the hexagonal closure architecture) is its non-circularity hinge — the candidate is non-circular only if that architecture is fixed independently of the generation count, which the paper flags as a demonstrable claim it does not here establish rather than as secured; its growth-law and sum-rule premises are open — so it *reduces* "why three generations?" to three checkable claims and attaches a falsifier (a fourth generation), rather than closing the question. This is the paper's nearest approach to falsifiable physics, and it is marked a candidate, not a result.

Epistemic markers: ***(inherited)*** *for results imported from prior VERSF papers;* ***(proven)*** *for results established here;* ***(conditional)*** *for results holding under stated inputs;* ***(conjectural)*** *for interpretive identifications;* ***(open)*** *for what remains undecided.*

Table of Contents

1. Introduction
2. Persistent Fold Defects and Representation Classes
    - 2.1 Persistent Fold Defects
    - 2.2 Representation Classes
3. The Particle Identity Problem and the Identification Premise
4. Conditions for Particle Identity
    - 4.1 The Restricted Refinement–Transport Bridge
    - 4.2 Why the Restriction Is Non-Circular Here
    - 4.3 The Representation-Class Persistence Principle
    - 4.4 Species versus Object
5. The Particle Species Identification Principle
    - 5.1 Sufficiency, Argued
    - 5.2 Necessity, Argued but Conditional
6. Indistinguishability — and What It Does Not Deliver 6A. Relation to Existing Accounts of Quantum Identity
    - 6A.1 Why Physicists Already Behave This Way
7. The Stability Spectrum
8. Gauge Bosons
9. Generations
    - 9.1 Candidate Generation-Count Derivation
    - 9.2 The Dyadic Loading Problem
    - 9.3 Depth Is Not a Direction 9A. Species Structure and the Standard Model 9B. Why Species Are Not Primitive 9C. Consequences for Open Standard Model Problems 9D. Why This Reduction Matters
10. Mass
11. Species Carrier Uniqueness

- 11.1 What Any Carrier of Particle Identity Must Satisfy
- 11.2 Representation Classes Satisfy All Four; the Alternatives Fail
- 11.2A The Collapse Dilemma: Even a Successful Carrier Becomes a Class
- 11.3 The Engine: No Surplus Identity
- 11.4 The Operational Species Criterion — and the Operational Species Principle
- 11.5 The Uniqueness Claim, at Honest Strength

12. Failure Modes
13. What the Paper Establishes 13A. What This Paper Does Not Reach

- 13A.1 What Would Convert This into Physical Support — a Roadmap

14. Conclusion

---

1. Introduction

The Standard Model begins with particles. Electrons, quarks, neutrinos, gauge bosons and the Higgs field are treated as primitive ingredients whose interactions are then specified.

VERSF reverses this order. The programme begins with committed distinctions and derives progressively richer structures. Previous papers established the hierarchy

Fold → Fact → Record → Persistent Fold Defect,

and subsequent work showed that Persistent Fold Defects organise into stable representation classes characterised by invariant tuples **(inherited)**.

The unresolved question is downstream of all of this:

Why should a stable representation class appear experimentally as a particle species, rather than as a merely abstract topological structure?

This paper addresses that identification problem, and its contribution is deliberately limited. It does **not** derive the Standard Model spectrum, exact masses, couplings, or exchange statistics. Its goal is narrower: to determine whether — and under exactly which conditions — particle *identity* can be reduced to stable representation classes of Persistent Fold Defects, and to mark precisely where that reduction is argued, where it is conditional, and where it remains open.

The discipline the paper holds itself to is the companion Object Criterion paper's: no inheritance is claimed that the prior results do not actually deliver, the persistence half of identity is carried only as far as the *restricted* Refinement–Transport Bridge allows, and the difference between a sufficient condition and a necessary one is kept visible throughout.

A note on genre, stated up front because it governs how the rest should be read. This is a paper of *conceptual analysis* — an ontological reduction — not a predictive physics paper. It contains no quantity that is forced to a value an experiment could confirm or refute: no mass, coupling, lifetime, mixing angle, generation count, or cross-section is derived here, and each is disclaimed explicitly where it might be expected. What the paper offers is an *identification* (particle species are representation classes of persistent closure structure) together with an argument that the identification is internally coherent, ontology-subtractive, and consistent with how physics already treats identical particles. That is a real contribution of the kind ontological reductions make — comparable in *form* to "temperature is mean molecular kinetic energy" or "the gene is a stretch of DNA" — but with an important caveat those analogies make vivid: *those* reductions earned their standing by going on to force measurable consequences (the gas laws, base-pairing) that could have failed and did not. This paper is at the prior, conceptual stage — the proposal of a reduction, not its experimental vindication — and it should be judged as such. The one place the wider programme could reach a genuine falsifier is the generation-count question: a derivation of why closure depth terminates at three, from an independently-fixed bound, would be refutable by a fourth generation. This paper *attempts* that derivation as an explicitly-marked **candidate** (§9.1) — and the candidate is instructive precisely about the line this note draws: its budget premise (a closure capacity K = 7 inherited from the hexagonal architecture) is non-circular *only if* that architecture is fixed independently of the generation count — a claim the paper flags as needing demonstration rather than asserting — and its growth-law premise (dyadic loading for arbitrary depth) is open, so the candidate *reduces* "why three?" to a small number of checkable claims rather than closing it. It is a candidate, not a result, and the paper marks it so. Until that premise is discharged generation-blind, this paper's claims remain conceptual, and the marker discipline below tracks *internal* provenance, not external confirmation, which the paper does not yet claim to have.

---

2. Persistent Fold Defects and Representation Classes

2.1 Persistent Fold Defects

A Persistent Fold Defect (PFD) is a refinement-stable closure structure satisfying:

1. non-trivial closure topology;
2. non-trivial closure holonomy;
3. refinement persistence;
4. positive closure stability.

The internal structure of a PFD is encoded by its invariant tuple

$\mathcal{I}(D) = (C\_D, \beta_1(D), h\_D, \pi\_D, \chi\_D, \gamma\_D, \ell\_D, \rho\_D)$,

whose entries capture closure completeness, first Betti number, holonomy, parity, chirality, generation depth, localisation, and confinement status respectively. The existence and refinement-stability of such structures is not established here; it is taken from the matter programme.

Status: (inherited).

2.2 Representation Classes

A representation class is the equivalence class

$[D] = \{ D' : \mathcal{I}(D') = \mathcal{I}(D) \}$.

All members of a representation class share the same admissible invariants, and therefore the class is preserved by any operation that preserves the invariant tuple. Representation classes accordingly survive admissible gauge transformations, admissible coarse-graining, and admissible refinement.

Status: (inherited).

One point of care, carried forward to §3: refinement-survival of a representation class is *description-invariance* — the class is the same under a change of resolution. It is **not** by itself endurance across the accumulation of fact. That distinction, established in the Object Criterion paper, is the whole subtlety of condition (4) in §4 and must not be smoothed over by the word "survive."

---

3. The Particle Identity Problem and the Identification Premise

A particle species is ordinarily treated as a physically identifiable entity. The question this paper poses is: *what structure actually carries that identity?*

Suppose two substrate structures possess identical admissible invariants. Exactly one of the following holds:

1. some admissible observable distinguishes them — in which case the invariant catalogue is *incomplete*, there being a physical difference the invariants fail to register; or
2. no admissible observable distinguishes them — in which case the distinction between the structures carries no operational content.

The paper adopts route (2), but does so *conditionally*, and is explicit that route (2) presupposes something route (1) names as a failure mode.

Identification Premise (conditional)

All admissible observables depend only on admissible invariant structure. Equivalently: the invariant catalogue $\mathcal{I}$ is *complete* — there is no admissible observable sensitive to a difference the invariant tuple does not record.

Status: (conditional); the completeness of the invariant catalogue is the specific thing assumed.

Two honesty conditions on this premise, neither of which an earlier framing made explicit.

First, non-circularity. If "admissible observable" were *defined* as "a functional of the invariant tuple," the premise would be true by definition and the indistinguishability it yields would be empty. The premise is contentful only if "admissible observable" has an independent characterisation — supplied by the distinguishability-geometry and Fact-Production sectors, where the admissible observables are the commitment-realisable distinctions, characterised *prior* to and independently of the invariant catalogue. The premise then asserts the substantive claim that those independently-characterised observables happen to factor through $\mathcal{I}$. That is the inherited content, and it is inherited from those sectors specifically, not from a generic "invariant-based ontology."

Second, catalogue-completeness is the live risk. Route (1) is not a logical curiosity; it is the precise way the premise can fail — a found admissible observable sensitive to a difference $\mathcal{I}$ omits would show the catalogue incomplete and reopen individuation. The premise is therefore conditional on catalogue-completeness, and we mark it so rather than presenting route (2) as forced.

Under the premise, members of the same representation class are operationally indistinguishable — a result used in §6 and there sized carefully.

---

4. Conditions for Particle Identity

The Object Criterion paper established that persistence alone is insufficient for objecthood: a structure can endure and still fail to be a thing if it is never individuated. The same caution applies to particle identity, and for the same reason.

We define a candidate particle species as a representation class satisfying:

1. refinement persistence;
2. closure stability;
3. representation invariance;
4. re-identifiability across admissible persistence intervals.

Conditions (1)–(3) are inherited directly from the PFD programme **(inherited)**. Condition (4) is the subtle one, and it does **not** follow from (1)–(3).

4.1 The Restricted Refinement–Transport Bridge

The Object Criterion companion paper showed that *refinement persistence* (invariance across a change of resolution) and *commitment-increment persistence* (endurance across the accumulation of fact) are relations along different axes, and that the unrestricted bridge between them —

invariance under refinement entails endurance across increments

— is not merely unproven but **false as stated**: an unstable particle is refinement-invariant (refine the lens and the defect is topologically still there) yet does not endure across increments — it decays. The unrestricted bridge is therefore rejected here as it was there.

What survives is the restricted form.

Restricted Refinement–Transport Bridge (conditional)

Re-identifiability across commitment increments holds over *disruption-free persistence intervals* — intervals in which no admissible decay or species-changing process occurs — and terminates when such a process does occur.

Status: (conditional); the open question is whether the *restricted* bridge holds, the unrestricted one being settled negative.

Consequently:

- an electron satisfies condition (4) over its observable lifetime;
- a muon satisfies condition (4) until it decays — its eventual decay does not undermine its identification as a species, it merely bounds the interval over which re-identifiability applies;
- the bridge is never assumed beyond the persistence interval itself.

This is the same restriction the companion paper reached, and it is inherited at that paper's status — re-identifiability for PFD-objects is conditional on the restricted bridge, not on an open general-transport account, which is a localised and tractable dependency rather than a blanket one.

4.2 Why the Restriction Is Non-Circular Here

The restricted bridge is in danger of vacuity in exactly one way, and this paper is positioned to disarm it where the companion paper could only flag it. The danger: if "disruption-free interval" means merely "an interval over which re-identifiability happens to hold," the restriction reads as *re-identifiability holds until it doesn't*, and asserts nothing.

The escape requires that "decay or species-changing process" be characterised **independently** of "the failure of re-identifiability" — and this paper has the machinery to do so in its own terms. A species-changing process is a *transition out of the representation class* [D]: a process after which the structure no longer satisfies $\mathcal{I}(D') = \mathcal{I}(D)$, so that it is no longer a member of [D]. Class membership is fixed by the invariant tuple (§2.2), entirely prior to and independent of any transport question. So the persistence interval is bounded by *class membership*, not by re-identifiability:

the disruption-free interval of a structure is the maximal interval over which it remains a member of its representation class [D].

This turns the restricted bridge from "holds until it fails" into "holds for as long as the structure remains in its class" — a contentful statement, because class membership is independently defined.

The boundary is empirical, not merely formal. Crucially, a class transition is not an abstract bookkeeping event: it is *observable independently of the transport question*, because the invariant tuple is built from admissible observables. The change is registered in the laboratory:

$\mu \to e + \bar{\nu}_e + \nu_\mu$: γ_D changes (generation 2 → 1), and the lepton-number bookkeeping changes — $\mathcal{J}$ takes a different value.

$n \to p + e + \bar{\nu}_e$: the closure/confinement and charge invariants change — a structure leaves the neutron class and a structure in the proton class is present.

In each case the invariant tuple *demonstrably* changes, and the change is measured directly — decay products, charge, generation — without any appeal to whether re-identifiability "held." So the interval boundary inherits the empirical content of the invariant catalogue: the moment a structure leaves [D] is the moment an observable invariant changes, which is exactly what decay experiments detect. The non-circularity is therefore not a definitional trick but a consequence of class membership being an observable matter of fact.

4.3 The Representation-Class Persistence Principle

The two preceding subsections together yield a clean statement of how species persistence relates to class membership. It is worth stating explicitly — but it must be stated for what it is, which is a *conceptual consequence of the identification adopted in this paper*, not an empirical theorem about nature. An earlier draft of this paper presented it as a proposition with a proof and called it the paper's distinctive result; that was an overstatement, and the honest framing is the weaker and more defensible one given here.

Representation-Class Persistence Principle (conceptual consequence of the identification; not an empirical theorem)

Once species identity is identified with representation-class membership (§4–5), persistence of species identity is *represented by* preservation of class membership: a structure persists as the species [D] over exactly the interval during which it remains a member of [D].

This is an unpacking, not a discovery, and the paper says so plainly. Under the identification, "the species persists" and "class membership is preserved" are two descriptions of one condition — the identity-defining invariants are unchanged. Spelling it out:

```
species persists  ≡  member of [D]  ≡  𝓙 preserved
```

All three are the same predicate in different vocabulary, because [D] is *defined* as the set of structures sharing $\mathcal{J}$ ($[D] = \{ D' : \mathcal{J}(D') = \mathcal{J}(D) \}$, §2.2), and species identity is *identified* with class membership (§5). So the "biconditional" an earlier draft proved — re-identifiable iff in-class — reduces, once both sides are unpacked, to *the identity-defining invariants are preserved iff the identity-defining invariants are preserved*. That is analytic. It is true, and it is useful as a clarification of what the identification commits us to, but it is not a fact about the world that could have come out otherwise. Presenting it with a proof and a ■ would dress a stipulation as a finding, and we do not.

What the principle *does* contribute is genuine but modest: it fixes the **vocabulary** in which species persistence is to be discussed. It says that, having made the identification, one need not treat "how long does this species persist?" as a separate transport question — it is answered by "how long does it remain in its class?", which is settled by the invariant tuple. The §4.2 result that the interval boundary is empirically observable (a class transition is a measurable invariant change) is the part with external content; the principle that persistence *tracks* class membership is the conceptual consequence that lets §4.2's observability do its work. The two are complementary: §4.2 supplies a non-circular, observable interval boundary; §4.3 notes that, under the identification, that boundary *is* the boundary of species persistence — by stipulation, not by theorem.

One reading to head off, because the equivalence invites it. "Persists iff in-class" concerns the *persistence of one individual member* of a class; it does **not** identify distinct members of the same class with one another. Two electrons share a class and are non-individuated *as to species* (§6), but they remain two — same class is not same individual. The principle says each electron persists as an electron for as long as *it* stays in the electron class; it says nothing that would collapse two class-members into one. The species/object distinction of §4.4 keeps these apart: the principle concerns an individual's persistence interval (an objecthood notion) indexed by a species-grade fact (its class), not a claim that class membership exhausts individuality.

Status: (conceptual consequence of the identification); analytic given the §2.2 class definition and the §5 species identification. It carries no empirical content of its own; the empirical content sits in §4.2 (the interval boundary is an observable invariant change) and is logically prior to, not derived from, this principle.

A note on honest strength: the principle does not establish that the restricted bridge holds (§4.1's open conditional), nor that the invariant catalogue is complete (§3's conditional). It says only that, *granting the identification*, species persistence and class membership are the same condition described two ways.

4.4 Species versus Object

Because this paper sits beside the Object Criterion paper, one distinction must be made explicit to prevent the two from being conflated. They answer different questions about the same structure.

- **Objecthood** asks: *what persistent individual structure is this?* It concerns a particular bundle of records — *this* atom, *this* defect — and is satisfied by re-identifiability plus approximate closure of that individual instance.
- **Species** asks: *which representation class does that structure belong to?* It concerns class identity — electron-kind, muon-kind — and is satisfied by the invariant tuple, which is shared across all individuals in the class.

```
Objecthood  →  individual instance  (this electron, persisting and closed)
Species     →  class identity        (electron-kind, fixed by 𝑱)
```

The relationship is that an object *of a given species* is an individual instance whose class membership is [D]. Objecthood individuates *within* the world of persistent structures (this one is not that one); species individuates *between* classes (electrons are not muons) while leaving members of one class non-individuated from each other (§6). The Object Criterion supplies the persistence-and-closure of the individual; this paper supplies the class identity of the kind. The Representation-Class Persistence Principle (§4.3) is the hinge between them — it notes that the individual's persistence interval (an objecthood notion) is described, under the identification, by its class membership (a species notion) — but as §4.3 stresses, that hinge is a conceptual consequence of the identification, not a derived bridge between the two papers.

---

5. The Particle Species Identification Principle

The principle is stated with its two directions explicitly separated, because they have different standing — the forward (sufficiency) direction is argued (§5.1), and the converse (necessity) direction is also argued, by the collapse dilemma of §11 (§11.2A), conditional on catalogue-completeness rather than left open.

Particle Species Identification Principle (conditional)

Under

1. fold-exhaustion,
2. the Identification Premise of §3 (catalogue-completeness),
3. the restricted Refinement–Transport Bridge of §4,

stable particle species *correspond to* stable PFD representation classes — where the forward direction (representation classes are *sufficient* carriers of particle identity) is argued in §5.1, and the converse (representation classes are *necessary* — nothing else could carry it) is argued in §11 via the collapse dilemma (§11.2A), which shows any carrier at all either fails observability or collapses into the class.

Note on the necessity direction: an earlier framing listed a separate fourth assumption — "no additional admissible carrier exists" (enumeration-completeness) — as an open premise the necessity direction rested on. The collapse dilemma of §11.2A removes that separate premise: it does not require the enumeration of alternatives to be complete, because *any* carrier, named or not, falls into its observable/unobservable split. So the necessity direction is conditional on catalogue-completeness (assumption 2) and not on a distinct enumeration-completeness premise.

The two have been consolidated; assumption (2) now carries both the forward and the converse direction's deepest dependency.

5.1 Sufficiency, Argued

Representation classes already provide, by inheritance:

- invariant structure (§2.2),
- refinement persistence and closure stability (§4 conditions 1–2),
- operational indistinguishability (§3, under the Identification Premise).

The remaining ingredient is re-identifiability across admissible persistence intervals (condition 4), supplied conditionally by the restricted bridge (§4.1) and made interval-contentful by class-boundedness (§4.2). Under the identification, this persistence is described by class membership (the Representation-Class Persistence Principle, §4.3) — a conceptual consequence of the identification, not an independent result, but enough to fix the vocabulary: a species persists over exactly the interval it remains in its class.

Therefore a representation class satisfying conditions (1)–(4) possesses *every presently identified requirement for particle identity*. This is the sufficiency direction: such a class is enough to carry particle identity. It is argued — modulo the stated conditionals — not merely asserted.

5.2 Necessity, Argued but Conditional

The converse — that *only* a representation class could carry particle identity — is **not** left as a bare open flag; it is argued in §11 (Species Carrier Uniqueness), and in its strongest form (the collapse dilemma, §11.2A) the argument does *not* rest on having enumerated every alternative: any carrier at all either fails the observability requirement or, being observable, collapses into an invariant and hence into the class. So the necessity direction has substantive support — not a shrug, and not merely an elimination hostage to its list, but a dilemma that closes the list. What it still rests on is catalogue-completeness (the dilemma's observable branch needs "observable $\Longrightarrow$ entry of $\mathcal{I}$"), which the paper carries as a premise and marks falsifiable (§12, mode 1). Accordingly the Identification Principle is *conditional rather than theorem-strength*: "correspond" should be read as "are sufficient carriers, and — given catalogue-completeness — the unique adequate carrier, since any rival either is undetectable or becomes a class," not as an unconditional biconditional. This mirrors the Object Criterion's care in distinguishing the sufficiency and necessity halves of its own biconditional, and inherits the same discipline — while improving on it, since the collapse dilemma gives necessity a footing the elimination alone could not.

---

6. Indistinguishability — and What It Does Not Deliver

The framework yields particle indistinguishability cleanly. Two electrons correspond to two instances of the same representation class; by the Identification Premise of §3, no admissible observable distinguishes them, so no admissible individuality remains between them. (The inference uses §3's premise directly and therefore inherits its conditional, catalogue-completeness status.)

This relocates to the level of PFD classes the identity that the Standard Model states at the level of irreducible representations: two electrons are identical there because they belong to the same representation, here because they instantiate the same PFD class. The explanation is structural rather than postulated.

The reach of this must be stated exactly, because it is easy to overclaim and a companion-paper draft did. What the identification delivers is *indistinguishability* — the statement that the configuration space of identical class-members is a *quotient*, with members not separately labelled. It does **not** deliver bosonic or fermionic *exchange statistics*. The boson/fermion dichotomy depends on the topology of that configuration space (the relevant fundamental group — the symmetric group in three dimensions, the braid group in two), on which one-dimensional representation is selected (the ± of symmetrisation versus antisymmetrisation), and on the spin-statistics connection — all of which carry geometric and relativistic content that no amount of non-individuation supplies. Mere class-identity tells you the space is a quotient; it tells you nothing about *which* exchange representation is realised or why that tracks spin.

Status: indistinguishability (the quotient structure) **(conjectural**, at the Identification Premise's conditional standing**)**; exchange statistics proper a **separate and larger open problem**, to be addressed by the spinorial and relativistic structures developed elsewhere in the programme, not here. The framing "reproduces exchange symmetry" is explicitly *not* claimed.

6A. Relation to Existing Accounts of Quantum Identity

Because this paper's subject — identity, individuality, and indistinguishability — is one with a substantial existing literature in the philosophy and foundations of physics, an outside reader is entitled to ask how the present account relates to that work, and what (if anything) the fold substrate adds rather than restates. This section locates the paper against the principal existing positions. It does not claim priority over them; in several places the VERSF account *recovers* a position already articulated elsewhere, and saying so is the honest way to show what is genuinely new versus what is a re-derivation in different vocabulary.

Ontic structural realism (Ladyman, French). The view that structure is ontologically prior to objects — that relations are fundamental and "objects" are derivative nodes — is closely allied to what this paper does: particle species emerge as representation-level structure, not as primitive individuals. The VERSF contribution here is not the structuralist thesis itself, which is prior, but a specific *substrate realisation* of it: the "structure" is the committed-fold record-network, and the derivative "objects" are persistent closure classes. Where generic structural realism asserts that objects reduce to structure, this account exhibits a candidate *mechanism* of the reduction. That is an addition of specificity, not of thesis; the structuralist will find the conclusion familiar and may reasonably ask whether the substrate adds explanatory power or only detail.

French & Krause; quasi-set theory. French and Krause argue that quantum particles are *non-individuals* — entities for which the usual identity relation does not meaningfully apply — and develop quasi-set theory as a formal framework in which cardinality without identity is coherent. The present account reaches a structurally similar conclusion (members of a PFD class are non-individuated, §6) but by a different route: not by denying identity its application, but by *locating* identity at the class level and showing there is no sub-class invariant to individuate members. The difference is worth stating precisely, because it could be a genuine divergence or merely a notational one: quasi-set theory denies that "same or different?" has a fact of the matter for the particles; this account says "same or different?" has a determinate answer at the level of *class membership* (same class) and no answer at the level of *which member* (no individuating invariant). Whether that is a real alternative to quasi-sets or a model *of* them is a question this paper raises and does not settle — and flagging it is more useful than claiming novelty.

Saunders' weak discernibility. Saunders, reviving a Quinean device, argues that fermions are *weakly discernible* — distinguished not by monadic properties but by irreflexive relations (no electron bears to itself the relation it bears to another). This is in tension with the present account's claim that same-class members are individuated by *nothing*: weak discernibility says there *is* a relational difference, the account here says there is none beyond class membership. The two need not conflict if "individuation" is read differently in each (Saunders individuates by relations, this paper by invariants), but the apparent disagreement is real enough that an honest comparison must flag it rather than paper over it: if weak discernibility succeeds, then same-class members are *relationally* discernible after all, and §6's "individuated by nothing" would need restriction to *monadic-invariant* individuation. The paper's claim should be read with that scope — non-individuation is by admissible invariant, and whether relational weak-discernibility cuts under that is not resolved here.

Quantum reconstruction programmes (Hardy; Chiribella–D'Ariano–Perinotti). These programmes derive the formal structure of quantum theory from operational or information-theoretic axioms (distinguishability, composition, purification). This paper's ambition rhymes with theirs — derive structure rather than postulate it — and its tools (admissible distinguishability, the Identification Premise that observables factor through invariants) are recognisably in the same family. The honest comparison is unflattering in one respect and clarifying in another. Unflattering: the reconstruction programmes *prove* that their axioms force the Hilbert-space formalism, the Born rule, and so on — they reach formal results that could have come out otherwise and didn't, which is exactly the external contact this paper lacks. Clarifying: this paper operates at a different and prior level — it is reconstructing *what a particle* *is*, ontologically, not *what the formalism is* — so it is not in competition with them but potentially upstream, supplying an ontology the reconstructed formalism could be interpreted over. Whether the fold substrate can be made to *force* a reconstruction-style result (rather than be consistent with one) is, again, the open question that separates a conceptual proposal from a derivation, and this paper is on the proposal side of that line.

What the comparison shows. Against this literature, the paper's genuinely distinctive element is narrow and should be stated as such: not the structuralist conclusion (shared with OSR), not non-individuality (shared with French & Krause), not the reconstructive ambition (shared with Hardy and CDP), but the specific claim that *a definite substrate — committed fold-records organised into persistent closure classes — realises all three at once*, with particle identity, indistinguishability, and the species/object distinction falling out of one invariant-tuple structure. That is a contribution of *unification and concreteness*, not of a new thesis about identity. It is worth having; it is not worth more than it is. And the comparison exposes the same gap §1 and §13A name from the other side: the neighbouring programmes that most resemble this one (Hardy, CDP) earn their standing by forcing formal results, and that is the step the VERSF account has not yet taken for particle ontology.

6A.1 Why physicists already behave this way

One point in the comparison is worth isolating, because it is the paper's *least* conditional claim and, for a physicist reader, perhaps its most persuasive. The identification proposed here is not a strange new criterion of particle identity imposed from outside. It is the criterion *already implicit in the practice of particle physics*, made explicit.

Consider how a physicist establishes that two objects are both electrons. There is no separate test of "electron-ness" beyond checking that they agree on a list of measured quantities:

```
same mass
same charge
same spin
same couplings
same lepton number
```

When all of these agree, the two are declared the same kind of particle — identical electrons — and no further question of "but are they *really* the same species?" is entertained, because there is nothing further to measure. That practice *is* the Operational Species Criterion (§11.4): two structures are the same species iff all admissible species-identifying observables agree. And the list a physicist checks is, item for item, the observable content of the invariant tuple $\mathcal{I}$ — mass and couplings downstream of the closure structure, charge and spin and lepton number among the invariants themselves.

So the paper is not proposing that physicists *should* identify species by invariant tuples. It is observing that they *already do* — that "same measured quantum numbers ⟹ same particle" is the working criterion of the field, and that the VERSF identification simply names what that criterion commits one to ontologically: if species identity *just is* agreement on the measured quantities, then a species *is* the class of structures sharing those quantities, which is the representation class. The contribution is to make explicit an identity criterion physics uses without stating, and to show it follows from the substrate rather than being an independent convention.

This claim is, notably, **less conditional than the rest of the paper**: it does not depend on fold-exhaustion or the restricted bridge. It depends only on the observation — uncontroversial — that particle physics individuates species by measured quantum numbers and admits no surplus criterion beyond them. That practice is itself a form of No-Surplus-Identity (§11.3) already operative in the field: physicists do not posit a fact about *which* electron is which beyond the quantum numbers, precisely because no measurement reaches such a fact. The VERSF account agrees, and explains why — the substrate provides no surplus for such a fact to consist in.

Status: (this paper; near-unconditional); an observation about existing practice, depending only on the uncontroversial fact that particle physics individuates species by measured quantum numbers, not on the programme's substrate premises.

Status: (this paper); a positioning against existing accounts, claiming unification and substrate-concreteness rather than priority, and explicitly flagging the points (quasi-set relation, weak-discernibility scope, reconstruction-style forcing) where the relationship is unresolved.

---

7. The Stability Spectrum

The Representation-Class Persistence Principle (§4.3) has an immediate consequence the earlier binary framing obscured: there is no sharp line between "stable" and "unstable" particles. Every particle species has a persistence interval — the interval over which it remains in its class — and species differ not in *whether* they persist but in *how long* the interval lasts. Stability is a spectrum, and all of it is species.

```
Stable class          →  electron, proton          →  interval unbounded (no
observed transition)
Metastable class      →  muon, neutron (free)      →  interval long but
finite
Rapidly decaying class →  short-lived resonances    →  interval very short
```

All three are genuine species in exactly the same sense: each is a representation class [D] fixed by an invariant tuple, each persists while it remains in its class (§4.3), and each undergoes a class transition — if it undergoes one at all — that is an observable change in $\mathcal{I}$ (§4.2). The differences are quantitative, in the length of the disruption-free interval, not categorical.

This matters for three reasons.

First, it dissolves a false dichotomy. The Standard Model treats a stable electron and a fleeting resonance as the same *kind* of entity — both "particles" — without a principled account of why one endures and one does not. Here the account is uniform: both are classes, both re-identifiable while in-class, differing only in transition rate. A resonance is not a defective particle or a mere "bump"; it is a representation class with a short persistence interval.

Second, it gives the framework empirical reach without overclaiming. The paper does not derive any lifetime (those are dynamical, downstream of the coupling and condensate structure of §10). What §4.2 does supply with empirical content is the *boundary*: a species transition is a measurable invariant change. The principle that persistence tracks class membership (§4.3) is, by contrast, a conceptual consequence of the identification, not an independent prediction — so the *falsifiable* part of the picture lives in §4.2 (a transition is an observable $\boldsymbol{\mathcal{J}}$-change), not in §4.3 (which is analytic). The honest claim is therefore narrow: *if* one accepts the identification, species transitions coincide with observable invariant changes, and that coincidence — not the analytic equivalence — is what could in principle be checked against decay data.

Third, it locates the truly stable species correctly. An electron and a proton sit at the top of the spectrum not because they are a different kind of thing but because no transition out of their class has been observed — their interval is, so far as measurement shows, unbounded. The framework does not *assert* their absolute stability (proton decay, were it found, would simply be a very-long-interval class transition); it represents stability as the limiting case of the same interval structure, which is the honest position.

Status: (this paper); a reading of the persistence picture under the identification (§4.3), with the *form* of the spectrum described and individual lifetimes left to the dynamical sectors **(out of scope here)**.

---

8. Gauge Bosons

The persistent transport programme identifies gauge transport with stable refinement-persistent cohomological structure **(inherited)**. Gauge-sector identity is therefore *provisionally inherited* through transport representation classes: a gauge boson species is a stable transport representation class, identified by the same machinery as a matter species, with interval-boundedness (§4.2–4.3) the one ingredient that remains open for the sector.

The openness is specific and worth naming rather than hedging. A propagating gauge excitation is typically massless and not localised, so the disruption-free interval — clean and class-bounded for a massive, localised defect like a muon — has no equally clean realisation for a virtual or in-flight gauge quantum, which has no lifetime-interval of the muon kind. The persistence principle (§4.3) applies *formally* (persistence described by class membership), but whether the class interval is well-behaved for non-localised transport structures is the open question. So gauge-sector identity is inherited at the same conditional standing as matter identity, *plus* one further open item — interval-boundedness — that the matter sector does not carry.

Status: (conditional); gauge identity provisionally inherited through transport representation classes, with interval-boundedness **(open)** for the sector.

---

9. Generations

The flavour programme introduces generation depth

$\gamma\_D \in \{1, 2, 3\}$

and the generation operator

$\mathcal{D} = \mathrm{diag}(1, 2, 4)$.

Generation is therefore a representation-class invariant rather than an independent ontological category. A muon and an electron differ because they occupy different representation classes — different admissible closure-depth sectors — not because they belong to different kinds of thing.

The muon's eventual decay does not undermine this identification. By the persistence principle (§4.3) it merely bounds the interval over which the muon-class structure persists as a muon: decay is a transition out of [D_muon] (γ_D: 2 → 1, plus the accompanying invariant changes), and the species persists precisely — and only — over the interval before that transition. Generation-as-invariant and decay-as-class-transition are two readings of the same class structure, consistent with each other.

Status: (inherited) for the generation structure; **(this paper)** for its consistency with the persistence account of §4.3.

A note on what the next three subsections are, and why they are here. Sections 9.1–9.3 push harder, and in a different register, than the rest of the paper — they reach toward a *falsifiable* consequence (a count of generations) where everything before §9 is conceptual reduction. This is deliberate, not a different paper smuggled in, and the reason is structural: the species reduction is precisely what *poses* the generation count as a well-defined question. Once a generation is a value of one representation-class invariant (γ_D) rather than a primitive particle family, "how many generations?" becomes "what is the admissible range of one invariant?" — a capacity question with a possibly-finite, possibly-derivable answer (§9.3). The ontology earns the physics question; the physics question could not even be posed sharply without the ontology. So §9 is where the conceptual reduction of §§1–8 cashes out into its one potential point of contact with data — the intended arc of the paper, ontology first *because* the ontology is what makes the generation question answerable. A reader who finishes thinking the generation material is the most alive part has read the paper correctly; a reader who concludes it is therefore a different paper has missed that §§1–8 are what license §9 to be posed at all. The candidate that follows is marked, throughout, as a candidate — its value is in *posing and localizing* the generation question, not in closing it.

9.1 Candidate Generation-Count Derivation

The paper has repeatedly named the generation count as the place the programme could first reach a genuine falsifier (§9C, §13A.1): the open question is *why* admissible generation depth terminates at three rather than continuing. This subsection states a candidate answer. It is offered as a **candidate theorem, not a result** — its conclusion follows from three premises, *none* of which is fully secured here: P1 (the closure budget) is conditional on a flavour-independence claim the paper flags but does not demonstrate, and P2 and P3 are open. The value of stating it is precisely that it *localizes* the open part onto three specific, checkable claims and attaches a clean falsifier, rather than leaving "why three?" diffuse. The honesty discipline of the rest of the paper applies here with extra force, because a generation-count derivation is exactly the kind of result a programme is tempted to reverse-engineer from the known answer — and the premise most at risk of that (P1, where 6+1 and 2^3-1 must coincide for unrelated reasons) is the one the paper is most careful not to call secured.

The candidate. Three generations are admissible because a generation-blind closure register of capacity K = 7 is exactly saturated by the first three dyadic closure-depth loads, and a fourth would overflow it:

```
generation depth loads (dyadic):   g₁ = 1,  g₂ = 2,  g₃ = 4,  g₄ = 8, …
closure register capacity:          K = 7
saturation condition (loads add):   g₁ + g₂ + ⋯ + gₙ = 2ⁿ − 1 ≤ K = 7

n = 3:   1 + 2 + 4 = 7   — saturates exactly
n = 4:   1 + 2 + 4 + 8 = 15 > 7   — overflows

⟹  exactly three admissible generation sectors.
```

Candidate Generation Capacity Theorem (candidate; conditional on three premises, two of them open)

Given (P1) a generation-blind closure register of capacity K = 7; (P2) dyadic generation loading, depth n carrying load 2^{n-1} for arbitrary n; and (P3) shared occupancy, so that admissible generation sectors jointly occupy the one register and their loads add — then exactly three generation sectors are admissible, since $2^3 - 1 = 7$ saturates the register while $2^4 - 1 = 15$ exceeds it.

The three premises must be separated and their standing stated exactly, because the theorem is non-circular only to the degree each is established *independently of the observed generation count*.

P1 — Closure capacity K = 7. (Conditional on demonstrated flavour-independence of the hexagonal architecture; this is the non-circularity hinge.) This is the premise that would most obviously make the argument circular if it failed, since $7 = 2^3 - 1$ is also "the sum of the first three dyadic loads" — so if K = 7 were derived anywhere downstream of the three-generation fact, the theorem would be 7 = 1 + 2 + 4 restated. K = 7 is inherited from the hexagonal closure architecture as the channel-count of a minimal complete closure cell — **six boundary channels plus one central hub** — and the programme presents it as fixed before generations or flavour enter. *If* that is so, the budget and the load-sum reach 7 by two independent routes (hexagonal channel-count on one side, dyadic saturation on the other), and their coincidence is the *content* of the theorem, not an assumption built into it.

But the independence must be *demonstrated*, not asserted, and the paper is obliged to flag that it is here asserted. The number 7 in the budget is 6 + 1, a hexagonal channel count; the number 7 in the saturation condition is $2^3 - 1$, a dyadic sum; the candidate's entire non-circularity is that these two sevens coincide *for unrelated reasons*. That holds only if the hexagonal cell — specifically, its *six* boundary channels — was selected by closure constraints that knew nothing about a downstream count of sectors. If anywhere in the matter or closure programme the hexagonal architecture was chosen, even partly, because it yields the right downstream multiplicities, then "6 + 1 = 7" is not generation-blind but generation-*informed through a longer cable*, and P1 is circular at one remove. The phrase "fixed before generations enter" asserts the priority; it does not show that the hexagonal architecture was fixed by constraints demonstrably independent of flavour. Everywhere else in this paper a hinge this load-bearing receives a demonstration or a downgrade; here, honestly, P1 has neither yet, so it takes the downgrade.

So P1 is **not** marked "secured." It is marked **conditional on a demonstrable fact the paper does not here establish**: that the hexagonal K = 7 architecture is fixed by closure constraints prior to and independent of the flavour/generation structure. Verifying that independence — tracing the selection of the hexagonal cell to constraints that could not have been tuned to the generation count — is itself a task, not a given, and it is the task on which the candidate's non-circularity actually rests. Until it is done, the coincidence of the two sevens is *plausibly* content but *possibly* construction, and the paper says so rather than resting the whole non-circularity claim on a parenthetical etymology.

What would close P1 — stated precisely, so the reader knows what success looks like, not only that the question is open. As with P2 and P3, the paper cannot perform the derivation here (it requires the closure architecture's formal content), but it can state exactly what a successful closure would consist of, so P1 is a defined task rather than a bare flag. The hinge is the question *why six boundary channels?* — and P1 closes if and only if the "six" is forced by a closure requirement that makes no reference to generations, flavour, or any downstream multiplicity. The candidate ground, to be confirmed or refuted in the closure programme, is this: **if six boundary channels are selected because closure requires a minimal complete cyclic transport cell — the smallest closed channel-cycle that admits a consistent central hub — then flavour never enters the derivation of six, and P1 is generation-blind.** On that route, "six" would be a fact about minimal cyclic closure (the smallest non-degenerate closed transport loop around a hub), fixed by transport-admissibility alone, with the generation count nowhere in its premises; the hexagon would be selected *as the minimal complete cyclic cell*, not *as the cell that yields three sectors*. The **refuter** is correspondingly definite: if anywhere the hexagonal cell is chosen because a different boundary-count would give the wrong downstream multiplicity — if "six" is justified by appeal to what it produces rather than by minimal cyclic closure — then P1 is generation-informed and circular at one remove. So the task is not vague: show that *six* is the minimal complete cyclic transport count, derived from closure/transport admissibility with no multiplicity input, and P1 is closed; find a multiplicity appeal anywhere in the selection of six, and it is refuted. That is what a successful P1 looks like, stated to the same precision as the closure conditions for P2 (the branching squeeze plus uniformity lemma) and P3 (the shared-channel-pool claim).

P2 — Dyadic loading 2^{n-1} for arbitrary n. (Open. The load-bearing gap.) This is where non-circularity actually turns, and at present it is *not* independently established. The programme has the generation operator $\mathcal{D}$ = diag(1, 2, 4) — but that is a *three-term* operator, and reading the doubling law off it presupposes three generations to write it down. For the theorem to derive three rather than assume it, the law "depth n carries load 2^{n-1}" must be established *for arbitrary n*, from the refinement structure itself — a recursion in which each refinement sector doubles the previous because of some closure-combinatorial fact, generating 1, 2, 4, 8, 16, … *before* the question of how many fit is asked. Such a derivation does not yet exist in the programme; $\mathcal{D}$ = diag(1,2,4) is inherited from the flavour-mixing programme as a three-term object, and the generation count is there marked conditional on it. So **P2 is the open premise**, and the candidate theorem's real effect is to *reduce* the question "why three generations?" to the sharper question "why dyadic loading for arbitrary depth, derived from refinement structure rather than fitted to three?" That reduction is the contribution; the theorem does not close the gap, it relocates it onto one precise, checkable claim.

P3 — Shared occupancy / the sum-rule. (Conditional; needs an argument the programme has not yet given.) Even granting P1 and P2, the saturation argument uses a third choice: that generation sectors *jointly occupy one register* so their loads *add* ($\Sigma \leq K$), rather than each sector merely having to fit individually (max-rule, $2^{n-1} \leq K$). The distinction matters for what the theorem explains. Under the *sum-rule*, three sectors saturate the register and a fourth overflows — this explains why *exactly three* exist. Under the *max-rule*, the constraint is only that each sector's own load fit under 7, which a fourth (load 8) fails — so the max-rule explains why a *fourth is too costly* but not why *exactly three* are present. The "exactly three saturate" result therefore needs shared occupancy specifically. That is a substantive claim — that generations are co-resident sectors of a single closure cell, jointly drawing on its seven channels — and while it is plausible (it is what "sectors of one cell" naturally means), the programme has not explicitly argued that the loads *add*. So P3 is conditional, pending that argument. (Note: both rules happen to forbid a fourth generation, which is reassuring for the *falsifier* below but means the arithmetic landing on $n \leq 3$ is weak evidence for the sum-rule specifically — P3 earns its place by the coexistence argument, not by giving the right answer.)

What would close P3 — stated precisely, so it is a task and not a hope. The paper cannot derive P3 here (it would require the formal content of the closure architecture — what a "channel" is, what it means for a sector to "occupy" one — which lives in the matter/closure programme, not this paper). But it can state exactly what must be shown, so the gap is a defined task. P3 closes if and only if the closure architecture establishes that **the seven channels of the register are a shared resource that distinct generation sectors draw on disjointly** — i.e. that a channel committed to one generation sector is thereby unavailable to another, so that the channel-occupancies of coexisting sectors partition (or at least sum within) the seven. Concretely, the claim to prove is: *the total channel-occupancy of the set of coexisting generation sectors equals the sum of their individual loads, and admissibility requires that total to be $\leq K$.* This is a statement about whether generation sectors share-and-compete for channels (sum-rule) or occupy channels independently/reusably (max-rule), and it is settled entirely by the closure architecture's account of channel allocation — an account this paper inherits but does not contain. The **refuter** is correspondingly definite: if the architecture permits two generation sectors to occupy the *same* channel without mutual exclusion (channels reusable across sectors), the loads do not add, the sum-rule fails, and the "exactly three" explanation collapses to the weaker max-rule "a fourth is too costly." So P3 is not vaguely "plausible"; it is the precise, checkable claim that generation sectors compete for a shared finite channel-pool, with reusable channels as its named falsifier — a task handed to the closure programme, with the success and failure conditions stated.

The falsifier, and why this is worth stating despite two open premises. The candidate's payoff is that it is *refutable*: it predicts no fourth generation, and a fourth generation — long searched for, so far absent — would refute it outright. That is exactly the contact-with-data the rest of the paper disclaims (§13A): unlike the conceptual reduction, this candidate could be wrong against experiment. It is the programme's most concrete falsifiable target, which is why §13A.1 names it first. And its structure is the right shape for a real derivation rather than a fit: an *independently-fixed budget* (P1 — conditional on that independence being demonstrated, §9.1) saturated by a *growth law* (P2), with the budget and the load-sum meeting at 7 by separate routes. What stands between candidate and theorem is precisely the three premises — P1's demonstrable flavour-independence, the dyadic law for arbitrary n (P2), and the shared-budget sum-rule (P3) — and naming them is the service this subsection performs: it tells the flavour and closure programmes exactly what must be proved, generation-blind, for three generations to be *derived* rather than inherited.

Status: (candidate; conditional). P1 **(conditional on demonstrated flavour-independence of the hexagonal architecture — the non-circularity hinge, asserted here, not yet shown)**; P2 **(open)** — the dyadic load law is not yet derived for arbitrary n, only read off the three-term $\mathcal{D}$; P3 **(conditional)** — the sum-rule needs a shared-occupancy argument the programme has not given. The conclusion (exactly three) is *not claimed as established*; what is claimed is that three follows from P1–P3, and that the derivation of three has thereby been reduced to three premises whose statuses are now explicit — with a fourth generation as the shared falsifier. The non-circularity of the whole candidate rests on P1's independence claim, which the paper flags as a task rather than a given (§9.1, P1). This is the candidate the roadmap of §13A.1 points to, stated at the only strength its premises currently support.

9.2 The Dyadic Loading Problem

Premise P2 of §9.1 — dyadic loading, depth n carrying load 2^{n-1} — is the one on which the candidate's non-circularity turns, so it deserves a section of its own examining what would discharge it. This section makes some progress and is candid that the progress is, once again, a *relocation* rather than a closure: it narrows P2 to a single sharper question, and identifies a substrate-level reason to expect the answer, without claiming to prove it.

The problem, stated precisely. The flavour programme supplies $\mathcal{D}$ = diag(1, 2, 4), but a three-term operator does not establish a general law. The entries 1, 2, 4 are consistent with infinitely many continuations — 1, 2, 4, 8, … (dyadic); 1, 2, 4, 6, … (then linear); 1, 2, 4, 4, … (saturating); and so on — and reading "doubling for all n" off the first three terms is extrapolation, not derivation. So the question is *not* "why are the observed loads 1, 2, 4?" (they are observed) but:

What substrate principle forces arbitrary refinement depth n to carry load 2^{n-1}, independently of how many generations are observed?

Only such a principle makes the §9.1 candidate independent of the generation count it is meant to explain.

Candidate route: binary refinement growth. Suppose generation depth measures the number of independent closure sectors that must be tracked to preserve distinguishability at a given refinement level. If each admissible refinement step splits every distinguishable sector into two distinguishable descendants, then

$N_1 = 1$, $N_{n+1} = 2 N_n$, giving $N_n = 2^{n-1} = 1, 2, 4, 8, \dots$

— the dyadic sequence, generated without reference to the observed generation count.

But this relocates the doubling; it does not derive it — and the section says so. Exactly as §9.1 reduced "why three?" to "why dyadic loading?", this route reduces "why dyadic loading?" to "why does each refinement step branch into *exactly two*?" The phrase "splits every sector into two" *is* the doubling law, transposed from the load to the branching factor: binary branching and load-doubling are the same claim. So the dyadic sequence is the unique consequence of *binary* growth — but binary growth is precisely what is in question, and "unique-given-binary" is not "derived." The candidate route is honest progress only because the relocated question is sharper and nearer the substrate than the one it replaces, not because it answers it. Two sub-assumptions are now visible and must be marked: the **base case** $N_1 = 1$ (depth 1 carries one sector — a stipulation, and the "first load is 1" datum of $\mathcal{D}$ in another guise), and the **branching factor** of exactly 2 (the real hinge).

Why 2 is nonetheless the natural answer — a substrate-level motivation, not a proof. The relocation is worth making because, unlike "why three generations?", "why branching factor 2?" has a candidate answer in the programme's own primitive. The Fold is the *minimal* distinction — a single binary cut, the smallest possible "this, not that." A refinement built from folds is, at each step, the addition of one further minimal distinction, which is intrinsically two-valued: a sector either is or is not further distinguished by the new fold. So binary branching is not an arbitrary choice of growth rate among many; it is what refinement-by-minimal-distinction *is* — each fold doubles because each fold is one binary cut. This is a motivation, not a derivation: it makes 2 the expected branching factor (where 3 or 6 would need a non-minimal primitive the substrate does not contain), and it ties P2 back to the fold itself rather than to $\mathcal{D}$. But turning "expected" into "forced" requires the squeeze below.

What would actually force 2: a two-sided squeeze, not an assumption. It would be circular to *assume* branching factor 2 (that is just P2 again). The non-circular target is to *derive* 2 by excluding every other rate from both sides:

- **From below:** any refinement law growing *slower* than doubling (branching < 2 on average — sub-binary) fails to preserve distinguishability, because some pair of sectors distinguishable before the refinement step would be merged or left undistinguished after it, contradicting that refinement preserves committed distinctions.
- **From above:** any refinement law growing *faster* than doubling (branching > 2 — super-binary) violates closure admissibility, because it would introduce more distinguishable descendants than a single minimal distinction can admissibly commit, requiring a non-minimal (hence forbidden) primitive distinction in one step.

If both bounds hold, binary growth is the *unique survivor* — forced from below by distinguishability-preservation and from above by closure-admissibility — and P2 becomes a derived property of refinement rather than a pattern inferred from the observed generations. This squeeze, *not* the bare assumption "each step doubles," is the actual thing to be established; the squeeze derives 2, whereas assuming binary branching merely restates it.

A hidden lemma the squeeze needs: uniform branching. Both bounds are stated here as the *shape* of the argument, not performed — and the from-below bound in particular hides a lemma that must be discharged before it closes. "Branching < 2 *on average* fails distinguishability" does not follow from "each step branches by < 2," because a *non-uniform* scheme — branching by 2 on some steps and by 1 on others, averaging below 2 — could preserve distinguishability at every step while still falling short of doubling overall. The from-below bound as stated tacitly assumes the branching factor is *uniform across depth*; without that, the bound leaks, since the merge-some-distinguishable-pair argument applies to a genuinely sub-binary *step*, not to a sub-binary *average* achieved by alternating full-binary and trivial steps. So the squeeze is not "distinguishability forbids slower, admissibility forbids faster" simpliciter; it is that *plus* a **uniform-branching lemma** — that admissible refinement branches at the *same* rate at every depth, so that "average < 2" implies "some step < 2." That lemma is plausible (it would follow if every refinement step is the same kind of operation — one minimal distinction — which is the fold-based picture of §9.2), but it is a third thing to prove, and the squeeze is closeable only with it. The paper flags this rather than let the two-line squeeze read as nearer to a proof than it is: between the squeeze and a real derivation of branching-factor-2 sits the uniformity lemma, **(open)**.

The regress, made explicit. The generation-count problem has now been narrowed twice, and stating the chain plainly is the honest summary of where the programme stands:

```
Why exactly three generations?
   ↓  (§9.1, given P1 K=7 + P3 sum-rule)
Why does generation depth carry loads 1, 2, 4, 8, … ?   (P2)
   ↓  (§9.2, binary-refinement route)
Why does admissible refinement branch by exactly 2?
   ↓  (the squeeze: distinguishability-preservation from below + uniform-
branching lemma;
       closure-admissibility from above)
[ open — but a question about the branching factor of refinement-by-minimal-
distinction,
  with the Fold's binary character as the substrate-level reason to expect 2,
  and a uniform-branching lemma still to discharge before the from-below
bound closes ]
```

Each step is a genuine reduction: the open question gets smaller, nearer the substrate, and more constrained. The present terminus — "does admissible refinement necessarily branch binarily?" — is open, but it is the *right* open question, because (i) it is answerable in principle by a distinguishability/admissibility squeeze internal to the programme, (ii) it has a clear expected answer motivated by the fold's minimality, and (iii) settling it affirmatively would discharge P2, leaving the §9.1 candidate resting on P3 (shared occupancy) alone — a single remaining obstacle between the closure architecture and a derivation of exactly three generations.

Status: (open; reduced). §9.2 does not derive dyadic loading; it reduces P2 to the binary-branching question, supplies a substrate-level motivation (the Fold as minimal binary distinction) for why the branching factor should be 2, and identifies the two-sided squeeze (distinguishability-preservation; closure-admissibility) that would *force* 2 rather than assume it — while flagging that the from-below half of the squeeze hides a uniform-branching lemma that must be discharged first. The branching question and the uniformity lemma are **(open)**; the reduction and motivation are **(this paper)**. P2 is not closed — but it is now a sharply-posed, substrate-near question with an expected answer and a known (if not yet complete) proof-strategy, which is the most this paper can honestly claim for it.

9.3 Depth Is Not a Direction

A clarification about what generation depth *is* belongs here, because §§9.1–9.2 invite a misreading that would quietly undo the whole candidate. The dyadic-load picture might suggest that successive generations occupy progressively deeper positions along some additional *direction* of the substrate — a fourth spatial-like axis indexed by generation. The programme rejects that reading, and the rejection is not a stylistic preference: it is what makes the generation count a *capacity* question (which can have a finite answer) rather than a *range* question (which cannot).

Depth is a hierarchy, not a coordinate. Following the inherited result *Depth Is Not a Direction* **(inherited)**, refinement depth is not a physical coordinate analogous to a spatial dimension. A spatial direction has metric structure, locality, propagation, and reversible transport; refinement depth has none of these. It indexes successive stages of admissible coarse-graining and descriptive resolution — a bookkeeping hierarchy, not a place one moves through. This paper applies that inherited distinction to the generation question specifically **(this paper)**.

Why the distinction does load-bearing work — it justifies P1's *form*. This is the section's real contribution, and it connects directly back to §9.1. The candidate theorem there assumed a *bounded budget* (P1, K = 7) without arguing why the constraint should be a budget at all rather than something open-ended. §9.3 supplies that reason. Consider the two readings:

- *If depth were a direction*, additional generations could always be introduced by extending farther along it. Nothing would bound them; no finiteness principle would follow, and asking "how many generations?" would be like asking "how many positions along the x-axis?" — a question with no finite answer. On this reading the §9.1 budget would be unmotivated.

- *If depth is an admissibility hierarchy*, further generations appear only if further independent refinement sectors remain *available*, and availability is set by a finite closure architecture. The question becomes "how many distinguishable refinement strata can the closure register support?" — a *capacity* question, which can have a finite answer.

So the choice between "direction" and "hierarchy" is exactly the choice between a question that cannot terminate and one that can. §9.1's budget premise (P1) is the right *kind* of constraint only on the hierarchy reading — and §9.3 is what licenses it. Without this section, a reader could fairly ask why generations should be capacity-bounded at all; with it, the capacity framing is the natural consequence of depth being an admissibility hierarchy rather than an axis.

The capacity reading of 1 + 2 + 4 = 7 — and its dependence on P3. On the hierarchy reading, each refinement depth contributes one independent binary admissibility channel, the distinguishability load grows 1, 2, 4, 8, … (under the §9.2 binary-branching hypothesis), and the bound

$$1 + 2 + 4 = 7$$

represents *exhaustion of the finite closure register* supplied by the K = 7 architecture — not exhaustion of geometric space. A fourth generation would require the next distinguishability sector, of weight 8; the obstruction is lack of remaining closure capacity, not lack of room. One honesty mark, so the vivid "exhaustion of the register" language does not over-reach: reading the bound as *summed* capacity-exhaustion (1 + 2 + 4 against one register) presupposes the sum-rule — that the sector loads *add* in a shared register — which is precisely the conditional premise P3 of §9.1. So §9.3's capacity reading inherits P3's conditional status: it is the correct interpretation *given* shared occupancy, and it does not by itself establish shared occupancy. What §9.3 secures is the weaker, unconditional point — that depth is a hierarchy, so the count is a capacity question with a possibly-finite answer — which holds regardless of whether the specific summed bound is right.

The reformulation. Generations are therefore not positions in depth. They are admissible refinement strata supported by a finite closure architecture, and the generation-count problem is a problem of *finite distinguishability* — how many independent distinguishable strata the closure register admits — not a problem of geometric extent. This aligns the candidate with the broader programme, in which finiteness arises from admissibility and closure constraints rather than from arbitrary truncation of an underlying continuum. It is also why the programme can hope for a *derivation* of three rather than a stipulation: a capacity has a definite value to compute, where a range does not.

Status: (this paper; interpretive, resting on an inherited distinction). The claim that depth is a hierarchy not a direction is **(inherited)** from *Depth Is Not a Direction*; its application to the generation count — that this makes the count a finite-capacity question and thereby justifies the *form* of §9.1's budget premise — is **(this paper)**. The summed-capacity reading of 1 + 2 + 4 = 7 inherits the conditional status of the sum-rule (P3); the unconditional content is only that depth's being a hierarchy makes a finite count *possible to derive* rather than open-ended.

9A. Species Structure and the Standard Model

The representation-class identification is not merely an ontological relabelling that leaves the physics untouched. The invariant tuple $\mathcal{I}(D)$ that fixes a species already *contains* the structural quantities from which the later Standard Model programmes build their results. Reading the tuple entry by entry:

Invariant	Standard Model role it carries
C_D	closure completeness — whether the defect is a complete, admissible closure
$\beta_1(D)$	first Betti number — the topological class of the closure
h_D	transport holonomy — the gauge-coupling structure
π_D	parity sector
χ_D	chirality — left/right-handedness
γ_D	generation depth — the flavour generation
ℓ_D	localisation / charge sector
ρ_D	confinement status — whether the species is confined (quark-like) or free (lepton-like)

The point is structural, not numerical: the particle-species reduction does not merely *say* that electrons are representation classes, it identifies the representation class with a tuple whose entries are *exactly* the quantities the flavour, gauge, and matter programmes take as their inputs. Chirality, generation, holonomy, charge, confinement — each appears in the programme's downstream sectors as a property of a particle, and each is, by the identification of this paper, an entry in the invariant tuple that *constitutes* the species.

So the reduction supplies the **substrate carrier** of those quantities. Where the flavour programme writes γ_D, it is referring to the generation entry of the species' tuple; where the gauge programme writes a transport holonomy, it is referring to h_D; and so on. This paper does not derive how those entries take their particular values — that is the work of the respective programmes — but it does establish *what kind of thing* they are: invariants of a persistent closure structure, not labels attached to a primitive particle. The species reduction is therefore the foundation the other programmes stand on, and §9A makes that dependency explicit rather than leaving it as a coincidence of notation.

Status: (this paper) for the identification of tuple-entries as the substrate carriers; **(inherited)** for what each entry physically encodes, and **(out of scope)** for the derivation of any entry's value.

9B. Why Species Are Not Primitive

The paper's recurring claim — that species are derived, not fundamental — has so far been carried by individual sections. It is worth stating once, explicitly, as a ladder, and showing that *each rung removes ontology* rather than merely renaming it.

```
Particle identity
      ↓   removes: the particle as a primitive bearer of identity
Representation class
      ↓   removes: the individual realisation as physically significant
Invariant tuple 𝕴
      ↓   removes: identity as anything beyond a list of invariants
Persistent Fold Defect
      ↓   removes: the defect as a primitive object (it is closure structure)
Closure structure
      ↓   (grounds out in the substrate primitives: folds + commitment)
```

Each downward step is an ontological *subtraction*, and the claim of this section is that no step smuggles a primitive back in:

- *Particle identity → representation class.* The particle is no longer a primitive entity carrying its own identity; identity is relocated to the class (§4, §5). What is removed: the particle as an irreducible bearer.
- *Representation class → invariant tuple.* The class is not an extra structure over and above the invariants; it *is* the set of structures sharing a tuple ($[D] = \{ D' : \mathcal{I}(D') = \mathcal{I}(D) \}$, §2.2). What is removed: the class as anything beyond invariant-identity.
- *Invariant tuple → Persistent Fold Defect.* The tuple is not free-floating data; it is the invariant content of a PFD (§2.1). What is removed: the invariants as primitive numbers — they are properties of a closure structure.
- *Persistent Fold Defect → closure structure.* The PFD is not a primitive object; it is refinement-stable closure topology (inherited, matter programme). What is removed: the defect as a thing — it is an organisation of committed records.
- *Closure structure → substrate.* Closure structure grounds out in folds and commitment, the programme's only primitives. No further ontology is introduced.

At no rung is a new primitive added; at every rung one is removed. This is a result *different in kind* from any partial derivation of a Standard Model number — neither stronger nor weaker, but answering a prior question. A derivation of, say, a CKM element would be a *quantitative* result, exposed to data, resting on assumptions; this is a *structural* result, not exposed to data, resting only on the identifications the paper has already argued. The two are not on one scale to be ranked: a quantitative derivation could be confirmed or refuted by measurement, which this reduction cannot be (§13A), and that is precisely why it is a different kind of claim and not a competitor to one. What the reduction shows, completely within its own kind, is that the entire apparatus of particle identity — species, classes, invariants — reduces without residue to persistent closure structure, monotone in ontology. Its value is the value of a clarified foundation, not of a confirmed prediction; §13A states the limit that makes this distinction matter, and §9B should be read as making the *ontological* claim only, with no implication that it outranks the physics the programme has not yet done.

Status: (this paper); conditional, as the whole identification is, on fold-exhaustion, the Identification Premise, and the restricted bridge — but adding no primitive at any rung is established given those, and is the paper's central ontological claim.

9C. Consequences for Open Standard Model Problems

The identification has a reach beyond its own statement: it reframes several open Standard Model problems as questions about representation-class structure. The following are **consequences, not derivations** — each has the form *if particle identity resides at the class level, then this problem becomes that question* — and the paper claims only the reframing, not the solution. Stating them is what makes the species reduction a *gateway* into the programme's harder problems rather than an isolated ontology note.

Exchange statistics. If particle identity resides at the representation-class level (§4–6), then exchange statistics cannot arise from primitive individuality — there is none to appeal to. They must arise instead from *geometric properties of class transport*: the topology of the configuration space of identical class-members and the holonomy of exchanging them. This relocates the open problem (§6) from "why are particles indistinguishable?" — answered, the quotient structure — to "what is the admissible exchange sector of a spinorial closure class, and is the antisymmetric sector the unique one?" That is a narrower, structural question about closure transport, and it is the natural next target the identification exposes. **(consequence; the question is open, the reframing is this paper's.)**

Generations. If γ_D is an invariant of the representation class (§9), then generation structure is not an external label attached to particles but a *property of species identity* — a closure-depth sector. The open problem thereby sharpens from "why do particles come in generations?" to "why does refinement persistence admit closure depths 1, 2, 3 and terminate there?" — and §9.1 takes this further than mere reframing, offering a **candidate derivation**: three dyadic closure-depth sectors (loads 1, 2, 4) exactly saturate a generation-blind closure register of capacity K = 7, while a fourth (load 8) would overflow it. The candidate's budget premise (K = 7, inherited from the hexagonal architecture) is non-circular only if that architecture is flavour-independent — a claim the paper flags as needing demonstration, not secured — and its growth-law premise (dyadic loading) is open, which §9.2 narrows once more, to whether admissible refinement branches by exactly 2; so the chain *reduces* the generation count to a small number of sharply-posed, substrate-near questions rather than closing it. This is the consequence with the most physical reach: it carries a falsifier (a fourth generation), which is why §13A.1 names it the programme's nearest contact with data. **(consequence + candidate; the count is reduced**

through §9.1's P2/P3 and §9.2's branching question, not derived; the reframing and reductions are this paper's.)

CKM / flavour mixing. If flavour mixing is transport *between* representation classes, then CKM structure is a statement about *admissible transitions between species classes* — which class-to-class transports are permitted, and with what amplitudes. The open problem moves from "why this mixing matrix?" to "what are the admissible inter-class transport amplitudes, and what fixes their seed?" The identification does not compute any amplitude — and the flavour programme's seed (the Cabibbo overlap) is, by its own account, imported rather than derived — but it recasts the mixing matrix as a transition structure on the space of species classes, which is the object the flavour programme actually manipulates. **(consequence; the amplitudes and their seed remain open in the flavour programme, the reframing is this paper's.)**

Notice what these three have in common, and what the section does *not* do. It does not derive exchange statistics, the generation count, or any CKM element — all three remain open, and two of them are explicitly conditional even within their own programmes. What it does is show that *each becomes a well-posed question about representation-class structure once the species identification is granted* — that the identification is not a terminus but the foundation on which the generation, flavour, and statistics programmes are built. The progression of the paper's final third is therefore:

```
What is a species?                          (§4–5)
      ↓
Why species are not primitive.              (§9B)
      ↓
How species carry SM structure.             (§9A)
      ↓
How species connect to generations.         (§9C)
      ↓
How species connect to flavour mixing.      (§9C)
      ↓
How species connect to exchange statistics. (§9C)
```

This is the honest reach of the result: the species reduction is a gateway into the Standard Model programme, supplying the substrate carrier (§9A) and the ontological foundation (§9B) on which the open problems (§9C) are posed — without itself claiming to solve any of them.

Status: (consequence / reframing); every item in §9C is a conditional of the form "if the identification holds, this problem becomes that question," importing no open result and asserting no derivation.

9D. Why This Reduction Matters

A skeptical reader, having granted the identification, may still ask the practical question: *why should it matter whether particles are representation classes rather than primitive entities?* The reduction is not bookkeeping; it has four concrete payoffs, each already established or reframed above and gathered here so the practical consequence is visible in one place.

1. It removes primitive particle ontology. The Standard Model posits electrons, quarks, neutrinos, and gauge bosons as fundamental ingredients — a list of primitives accepted without explanation. The reduction eliminates that list: every entry on it becomes a stable representation class of persistent closure structure, with no primitive added at any rung (§9B). A theory with fewer primitives, all else equal, explains more; this reduction removes an entire ontological category.

2. It unifies matter and gauge sectors under one classification. Matter particles and gauge bosons are treated, in the Standard Model, as different kinds of object — fields of different spin and role. Here both are representation classes: matter species are stable defect-closure classes, gauge species are stable transport-representation classes (§8), but *both are classes of persistent substrate organisation*, identified by the same machinery. The matter/gauge distinction becomes structural (which sector of closure) rather than ontological (which kind of primitive).

3. It localises three open Standard Model questions onto specific invariants. This is the most practical payoff for the rest of the programme. Each open problem of §9C is now pinned to a *named entry of the invariant tuple*, turning a diffuse "why is the Standard Model like this?" into three sharp questions about three specific invariants:

```
exchange statistics  →  h_D + spinorial (double-cover) closure sector
                          "which exchange holonomy sector is admissible?"
chirality            →  χ_D
                          "why this parity/handedness assignment?"
generations          →  γ_D
                          "why do admissible closure depths terminate at 3?"
flavour transport    →  inter-class transport (the h_D-carried transitions)
                          "which class-to-class transitions are admissible?"
```

(Note the distinction kept here: exchange *statistics* is carried by the spinorial/exchange-holonomy structure — h_D and the double cover — while *chirality* is the separate invariant χ_D. The two are listed apart deliberately; conflating them would misstate which invariant each question lands on.) A programme that knows *which invariant* an open question concerns can attack it directly, rather than treating each particle property as an independent mystery.

4. It converts disparate particle properties into properties of one object. Spin behaviour, generation, charge, confinement, chirality, mixing — in the Standard Model these are a heterogeneous collection of attributes attached to particles. Under the reduction they are all entries of a single underlying object, the invariant tuple $\mathcal{I}$(D) (§9A). What looked like a list of independent particle properties becomes the structured content of one closure invariant. That is both an ontological economy and a research strategy: the questions are no longer "why does the electron have property X, and separately property Y?" but "what fixes the tuple $\mathcal{I}$ of the electron class?" — one question about one object.

The payoff, in a sentence: the reduction trades a long list of primitive particles with heterogeneous attributes for a single classification scheme in which every particle property is an entry of one closure invariant, and in which three of the Standard Model's deepest open questions become well-posed questions about three named entries of that invariant. That is why it matters — not because it relabels, but because it concentrates.

Status: (this paper); a gathering of the payoffs of §9A–9C, asserting no new result beyond the economy and localisation the reduction already supplies.

---

10. Mass

Mass enters only after identity is established. The closure-condensate framework gives

$m_D = g_D \cdot v$,

where v is the closure condensate and g_D is determined by the representation-class invariants **(inherited)**.

Mass is therefore a *derived* property of a representation class interacting with the condensate, not a constituent of particle identity: identity is fixed by the invariant tuple before any coupling to v is considered. Identity precedes mass.

This also supplies, in passing, part of the §4.2–4.3 machinery: the coupling g_D and the condensate dynamics are among the channels by which a class-transition (decay) can occur, so the "species-changing process" that bounds the persistence interval is characterised by the same inherited dynamics that fix mass — reinforcing that the interval boundary is independently specified, not circular.

Status: (inherited).

---

11. Species Carrier Uniqueness

This is the paper's central argument, and the place where its claim is most ambitious: not merely that representation classes *can* carry particle identity, but that — under the programme's ontology and two marked completeness conditions — they are the *only* available carrier. The earlier sections supply the pieces; this section assembles them into the uniqueness claim and states, with care, exactly how strong it is.

11.1 What any carrier of particle identity must satisfy

Suppose particle identity exists at all — that there is a fact about what makes a structure the kind of thing it is, persisting and re-identifiable. Then whatever *carries* that identity must meet four requirements. Each has a *role-level* statement — neutral, in terms any account of particle identity would grant — and a *programme-level* statement, which is the substrate's realisation of the neutral version. The distinction matters, and we keep it explicit, because two of the four are easy to state in a way that pre-shapes the result toward representation classes.

1. **Persistence** — *(role)* identity must endure: a particle remains itself as the world goes on. *(programme)* the carrier endures across commitment increments (§4). This requirement is role-forced and neutral: any account agrees a particle that did not persist would carry no identity.
2. **Observability** — *(role)* identity must be detectable: an identity no measurement could reach does no physical work and makes no physical difference. *(programme)* the carrier is accessible to admissible observation (§11.4). Also role-forced and neutral — and it is the requirement that does the real eliminative work in §11.2.
3. **Transportability** — *(role)* the same particle is the same particle under a change of reference frame; identity does not depend on who describes it. *(programme)* the carrier survives admissible transport and coarse-graining (§2.2).
4. **Refinement-invariance** — *(role)* the same particle is the same when examined more closely; identity does not evaporate under finer description. *(programme)* the carrier is stable under refinement (§2.1).

A note on (3) and (4), because a sceptic is owed it. Stated at *programme* level — "survives admissible coarse-graining," "stable under refinement" — these are close to "behaves like a representation class," since §2.2 *defines* representation classes to survive exactly those operations; a primitive-substance theorist could fairly object that their carrier need not survive coarse-graining in the programme's technical sense. So the requirements are not neutral *as stated in programme vocabulary*. The defence is that their *role-level* versions — frame-independence and stability-under-closer-examination — are neutral: any account of particle identity must grant that an electron is the same electron in another frame and the same electron under a microscope, on pain of the "identity" tracking the observer rather than the particle. The programme's admissible-transport and refinement operations are then offered as the *substrate's realisation* of those neutral requirements — the claim being that, in a world made of folds and records, "same under change of frame" and "same under finer description" *just are* survival under admissible transport and refinement. That claim is itself substantive (it asserts the programme operations correctly capture the neutral requirements) and is inherited from §2.1–2.2 rather than argued here. So (3) and (4) are role-neutral at the level of what they demand, programme-specific at the level of how the demand is cashed out, and the cashing-out is an inherited identification, not a neutral test. We flag this rather than let "read off the role" carry weight it cannot.

With that qualification, the four are read off the role at the neutral level: any candidate carrier — representation class or otherwise — must persist, be detectable, be frame-independent, and survive closer examination. Requirements (1) and (2) are neutral outright; (3) and (4) are neutral at role level and realised in programme terms by an inherited identification. The elimination in §11.2 turns mainly on (2), observability — the requirement least open to the "stated to fit classes" objection — which is worth noting, since it means the load-bearing strike does not depend on the (3)/(4) qualification.

11.2 Representation classes satisfy all four; the alternatives fail a named requirement

Representation classes meet the four directly: persistence (§4, conditional on the restricted bridge), observability (the invariant tuple is built from admissible observables, §11.4), transportability and refinement-invariance (both inherited PFD-class properties, §2.2). So a representation class is an *adequate* carrier.

The substantive claim is that the alternatives are not — each fails one of the four, or a premise the programme has adopted. Enumerate the candidates the vocabulary of "particle identity" suggests, and strike each against a *named* requirement or principle rather than a generic appeal:

Candidate carrier	Fails on	Specific ground
Primitive particle substance	requirement (2), and fold-exhaustion	An enduring substance beneath the records is not a substrate ingredient the programme admits (fold-exhaustion); and insofar as it is posited *beyond* the invariants, it is unobservable — it fails observability.
Haecceity / primitive thisness	requirement (2), admissibility	A bare individuating label distinguishing otherwise-identical structures registers in no admissible observable; under the Identification Premise (§3) it marks no physical difference. Fails observability.
Hidden labels beneath the invariants	catalogue-completeness (§3)	A difference no admissible observable records is, under the premise, no difference at all. Asserting one *is* the denial of catalogue-completeness, not an alternative carrier compatible with it.
Non-representational substrate markers	fold-exhaustion + catalogue-completeness	A marker not expressible through the invariant structure is neither a substrate ingredient the programme admits nor an admissible observable. Fails on both.

Each candidate falls to a *named* requirement or principle already load-bearing in the paper. That is what makes the elimination checkable rather than rhetorical: a reader who accepts the four requirements and the three principles (fold-exhaustion, the Identification Premise, catalogue-completeness) must accept the strikes; a reader who rejects a strike is thereby rejecting one of those named commitments, not the table.

But an elimination is only as strong as the completeness of its list — and the next subsection shows the list does not, in fact, need to be complete, because there is a dilemma that closes it.

11.2A The collapse dilemma: even a successful carrier becomes a class

The elimination table is hostage to a worry the companion paper also carries: maybe some carrier no one has named escapes every strike. The following argument retires that worry — not by enumerating better, but by showing that *any* carrier whatsoever, named or unnamed, either fails the observability requirement or collapses into the representation-class picture. There is no third option, so the enumeration need not be complete.

Collapse Dilemma (conditional on catalogue-completeness)

Let X be *any* proposed carrier of particle identity. Either X is admissibly observable, or it is not.

- **If X is not observable**, it fails requirement (2): an identity no admissible measurement can reach makes no physical difference (§3) and does no physical work. X is struck.
- **If X is observable**, then by catalogue-completeness (§3) X is registered by some admissible observable, hence is an entry of — or a function of — the invariant tuple $\mathcal{I}$. Then "same X" is "same on that invariant," and "same carrier" reduces to "same $\mathcal{I}$," hence to same representation class. X has not *rivalled* the class picture; it has *become* it.

Either way, X is not an alternative to the representation class: unobservable carriers are struck, observable ones collapse into the class. The dilemma is exhaustive by excluded middle on "observable or not," so it holds for carriers not yet imagined.

This is strictly stronger than the elimination table, and in a specific way: it removes **enumeration-completeness** as the load-bearing open premise. The table left uniqueness hostage to "have we named every candidate?"; the dilemma answers "it does not matter — any candidate, named or not, lands in one of two branches, and neither yields a rival." What the dilemma rests on instead is **catalogue-completeness** (the right branch needs "observable $\Longrightarrow$ entry of $\mathcal{I}$"), which the paper already carries as a premise. So the collapse argument *consolidates* two open premises into one: where §11.2's elimination needed both enumeration-completeness and catalogue-completeness, the dilemma needs only catalogue-completeness. That is a real reduction in what the uniqueness claim is hostage to — not the elimination of conditionality (catalogue-completeness remains a substantive, falsifiable bet, §12 mode 1), but its *consolidation* onto the single premise the paper most wants to stake itself on.

The table (§11.2) is retained, because exhibiting the named carriers and showing exactly how each lands in the dilemma's branches is more convincing than the abstract dilemma alone. But the table is now illustrative of the dilemma rather than load-bearing: even if it missed a candidate, the dilemma catches it.

11.3 The engine: No Surplus Identity

The reason every alternative either fails observability or collapses into a class is one underlying fact, worth stating on its own because it is what *drives* both the elimination and the dilemma — though it must be stated for what it is: a consequence of catalogue-completeness, not an independent discovery.

No-Surplus-Identity Principle (consequence of catalogue-completeness; not an independent theorem)

If catalogue-completeness holds (the admissible observables factor through the invariant tuple $\mathcal{I}$, §3), then there is no admissible physical difference between two structures beyond a difference in $\mathcal{I}$. There is no *surplus identity* — no residue of thisness, no hidden label, no sub-invariant fact — for an alternative carrier to consist in.

This is an unpacking of catalogue-completeness, and the paper does not dress it as more: catalogue-completeness *says* observables factor through $\mathcal{I}$, and "no surplus beyond $\mathcal{I}$" is that statement read from the identity side. It earns no ■. But stating it makes both arguments' structure visible. It is why the elimination's strikes are not accidents (each named alternative is an attempt to locate identity in a surplus that, under completeness, does not exist), and it is the *content* of the collapse dilemma's right branch (an observable carrier has no surplus to consist in beyond $\mathcal{I}$, so it collapses into $\mathcal{I}$). Primitive substance, haecceity, hidden labels are all names for surplus identity; No-Surplus-Identity says there is none; so all of them — and any unnamed candidate — fail or collapse, for one reason. The principle is the engine; the table is one output, the dilemma the other.

11.4 The Operational Species Criterion — and the Operational Species Principle

The observability requirement (2) is not abstract. It corresponds to what experiments actually do, and stating it operationally is what turns the uniqueness argument from metaphysics into something with a measurement procedure attached.

Operational Species Criterion (conditional on catalogue-completeness)

Two structures belong to the same species if and only if all admissible species-identifying observables agree on them. Under catalogue-completeness, this is equivalent to identity of the invariant tuple, hence to membership of the same representation class:

same admissible observables ⇔ same $\mathcal{I}$ ⇔ same class.

Both equivalences in this chain are **analytic given their premises**, and the criterion's content lives entirely in one premise being a substantive empirical bet — not in the chain being more than an unpacking. This is the same discipline applied to the §4.3 persistence principle, and it applies here for the same reason. Trace the left equivalence (same observables ⇔ same $\mathcal{I}$): catalogue-completeness *is* the statement that the admissible observables factor through $\mathcal{I}$ — that they do not reach past it. So "agreement on all admissible observables forces agreement on $\mathcal{I}$" is true *because completeness says observables don't reach past* $\mathcal{I}$; the equivalence is completeness restated from the identity side, exactly as §4.3 was the class definition restated. It is therefore analytic given completeness, and it inherits completeness's standing precisely — no more. The right equivalence (same $\mathcal{I}$ ⇔ same class) is the §2.2 definition, and is analytic outright.

So the honest description is *not* "substantive because completeness is assumed" — conditioning on completeness does not upgrade an equivalence's logical status, it transmits the premise's status to it. The correct statement is: the criterion is **analytic given catalogue-completeness, and inherits completeness's conditional, empirically-exposed standing**. Where the real content sits is in completeness itself, which is a non-trivial, falsifiable premise: a found sub-invariant observable would refute it (§12, mode 1). The criterion is "contentful" only in the derived sense that *the premise it unpacks could fail* — if completeness is false, "same observables" forces only "same observable-part-of-$\mathfrak{I}$," the chain breaks, and the species/class identification breaks with it. That is a genuine exposure to the world, but it is completeness's exposure, transmitted through an analytic equivalence, not content the equivalence generates on its own. We mark it so, rather than letting "not merely definitional" suggest the conditioning does logical work it does not do.

What the criterion buys, then, is not a stronger logical status but *concreteness with a located vulnerability*: it says every experimental procedure that identifies a particle's species — measure its charge, spin, generation, chirality, confinement — is, under completeness, *already a measurement of the invariant tuple* $\mathfrak{I}$, with no separate "species measurement" over and above tuple measurement; and it shows that this whole picture stands or falls with the single empirical premise that the catalogue is complete. This is taken up in §6A.1's "why physicists already behave this way," where the same identification appears as existing practice — and notably, the practice version is *less* conditional, because it rests only on the observed fact that physics admits no surplus criterion, not on a formal completeness premise.

The Operational Species Principle — a result of physics, not of VERSF

The previous paragraph contains the paper's most important strategic point, and it deserves to be lifted out of the VERSF setting entirely, because in its strongest form *it does not mention VERSF at all*. The reduction of species to equivalence classes is forced not by this programme's ontology but by the working practice of particle physics. State it as a principle belonging to physics generally:

Operational Species Principle (physics-general). Suppose: (i) species are identified experimentally; (ii) species-identification proceeds entirely by measurable properties; (iii) no admissible species-identifying criterion exists beyond the measurable properties. Then a species *is* an equivalence class of structures under agreement on the measurable identifying properties:

species = equivalence class of invariant observables.

It is called a Principle, not a Theorem, deliberately. The result is near-analytic given premise (iii) — premise (iii) is catalogue-completeness in physics' own clothes, and the conclusion unpacks it — exactly as the Operational Species Criterion is analytic given completeness (§11.4) and the No-Surplus-Identity Principle is an unpacking of it (§11.3). A hostile philosopher would rightly point out that an analytic-given-its-premises statement dressed as a "theorem" overstates its logical status; the collapse dilemma (§11.2A) can carry theorem-style language because it performs a genuine exhaustive case split, but this result cannot, because it unpacks a premise rather than proving something the premise did not already contain. So it takes the same epistemic label as No-Surplus-Identity: a *Principle*, near-analytic given its stated premise.

What gives it force is not its logical status but *whose* premise (iii) is. It is not a VERSF posit. It is the working assumption of particle physics: the field individuates electrons by mass, charge, spin, and couplings, and entertains no further fact about "which electron is which" beyond agreement on those numbers (§6A.1). Premises (i) and (ii) are simply descriptions of experimental practice. So the principle says: *granting the way physics already identifies particles, a species is an equivalence class of invariant observables* — a conclusion in the philosophy of physics generally, owing nothing to folds, records, or PFDs.

This re-layers the whole paper, and the re-layering is the strongest framing available to it:

```
Operational Species Principle       (physics-general; species = class of
invariant observables)
        ↓   forced by particle-physics practice, not by VERSF
representation class               (the equivalence class the principle
delivers)
        ↓   VERSF's distinctive, separable, more-contestable step
PFD representation class           (VERSF identifies the class with a
Persistent Fold Defect class)
```

The first arrow is owed to physics, not to this programme; the second arrow is VERSF's own contribution and carries all of VERSF's conditionality (fold-exhaustion, the matter programme, the rest). This split is worth making loudly because it *quarantines* the controversial part. A reader sceptical of the fold ontology can reject the second arrow entirely and still owe the first: species are equivalence classes of invariant observables *whether or not VERSF is right*, because that follows from how physics identifies particles. VERSF's claim is only the further, separable identification of *what the class is made of* — persistent closure structure. The general principle stands without the programme; the programme adds the substrate realisation.

Status: the Operational Species Principle is **(physics-general; near-analytic given premise (iii), which is existing practice)** — a Principle rather than a Theorem precisely because it unpacks its premise rather than exceeding it; its standing does not depend on the VERSF substrate. The PFD-identification (second arrow) is **(this programme; conditional on fold-exhaustion and the matter programme)** — the separable, more-contestable step. Keeping the two apart is what lets the paper claim the first without staking it on the second.

11.5 The uniqueness claim, at honest strength

Assembling §§11.1–11.4:

Species Carrier Uniqueness Argument (conditional on fold-exhaustion and catalogue-completeness)

Given fold-exhaustion and catalogue-completeness, the representation class is the unique adequate carrier of particle identity. Two arguments converge on this: the *elimination* (§11.2), which strikes each named alternative against a carrier requirement; and the stronger *collapse dilemma* (§11.2A), which shows that any carrier whatsoever — named or not — either fails observability (and is struck) or is observable (and collapses into an invariant, hence into the class). Both run on the No-Surplus-Identity engine (§11.3): under completeness there is no surplus beyond $\mathcal{I}$ for an alternative to consist in.

The strength of this claim must be stated exactly, because uniqueness invites overstatement and an earlier instinct of this programme was to give in to it. What it is, and three things it is **not**:

- It **rests on catalogue-completeness, and now on that alone** for the carrier question (fold-exhaustion still underwrites the substrate framing). This is the §11.2A improvement: the elimination of §11.2 was hostage to *enumeration-completeness* ("have we named every candidate?"); the collapse dilemma removes that hostage by the exhaustive observable/unobservable split, leaving only catalogue-completeness. The two open premises are consolidated into one — the one the paper most wants to stake itself on.
- It is **not** uniqueness *simpliciter*. It is uniqueness conditional on catalogue-completeness: if completeness fails — a sub-invariant observable exists — then an observable carrier need *not* collapse into $\mathcal{I}$ (it could consist in the sub-invariant difference), and the dilemma's right branch breaks. Catalogue-completeness is the deepest condition and the one most exposed (§12, mode 1); the uniqueness is exactly as strong as it.
- It is **not** the claim that no other carrier is *expressible*. The collapse dilemma does not say rivals cannot be formulated; it says any formulated rival either is undetectable (struck) or collapses into the class. That is a claim about where carriers *land*, not about what can be *said* — and it is the honest, defensible form, not the "inexpressibility" claim the companion paper withdrew.

With those scopes marked, the claim stands at its real strength: *given the programme's substrate framing and catalogue-completeness, the representation class is the unique adequate carrier of particle identity — not merely because every named alternative fails, but because any carrier at all either fails observability or becomes a class, there being no surplus identity for it to be made of.* That is a uniqueness claim about the paper's actual subject — sharper than "here is a theory of species," sharper even than the elimination alone (it no longer needs the enumeration to close), and weaker than "here is the only conceivable theory." It is exactly as strong as catalogue-completeness.

Status: (conditional on fold-exhaustion and catalogue-completeness); the collapse dilemma (§11.2A) retires enumeration-completeness as a separate load-bearing premise, consolidating the carrier question onto catalogue-completeness alone. The forward (sufficiency) direction (§5.1) is unconditional on either.

---

12. Failure Modes

The identification earns its standing as a *physical* claim, rather than a definition, by specifying exactly what would defeat it. There are four such conditions, each tied to a named premise, and each in principle checkable.

#	Failure mode	What it would break	The premise it bears on
1	A **sub-invariant observable** exists — an admissible observable sensitive to a difference $\mathcal{I}$ does not record.	Indistinguishability (§6) and the whole identity-at-class-level claim: class-members would become distinguishable.	The **Identification Premise / catalogue-completeness** (§3). *Empirical:* a found sub-invariant observable would refute it.
2	A **species transition is observed with no invariant change** — a structure changes kind while every admissible invariant $\mathcal{I}$ stays fixed.	The *usefulness* of the identification, not its consistency: it would show $\mathcal{I}$ fails to capture what "species" tracks, i.e. the identification was the wrong one — a sub-invariant observable in disguise (collapses to mode 1).	The **identification** (§5) and catalogue-completeness (§3). *Empirical, via mode 1.*
3	The **invariant catalogue is observed to be incomplete** — any case where two structures agree on $\mathcal{I}$ but differ observably.	Both the indistinguishability claim and the identification: the equivalence species ≡ $\mathcal{I}$ would no longer hold.	**Catalogue-completeness** (§3). *Empirical.* (Note: the §4.3 principle itself cannot be empirically broken — it is analytic given the identification; what is empirically at risk is the *identification*, mode 2/3, not the principle that unpacks it.)
4	**Another admissible carrier of identity is discovered** — a carrier escaping both the elimination and the collapse dilemma.	*Largely absorbed by the collapse dilemma (§11.2A).* Any discovered carrier is either unobservable (struck on observability) or observable (and collapses into an invariant, hence the class) — so a genuine *rival* requires an observable difference that $\mathcal{I}$ fails to record, which **is** mode 1. After §11.2A, mode 4 is not an independent failure mode but a special case of mode 1: the only way a new carrier rivals the class picture is if catalogue-completeness fails.	**Catalogue-completeness** (§3) — *via* the collapse dilemma; no longer a separate enumeration-completeness premise.

The asymmetry the recast makes explicit: the empirically vulnerable claims are the *identification* and *catalogue-completeness* (modes 1–3), not the §4.3 persistence principle, which — being analytic given the identification — cannot itself be falsified by any observation. An earlier draft listed "class transition without invariant change" and "re-identifiability surviving invariant change" as falsifiers *of the proposition*; that was a category error, since the proposition is a definitional consequence. What those observations would actually impugn is the *choice of identification* (mode 2) or the *completeness of the catalogue* (mode 3) — the premises with empirical content — and the table now attributes them correctly. Mode (4), after the collapse dilemma (§11.2A), is no longer an independent failure mode: a discovered rival carrier can only rival the class picture by exhibiting an observable difference $\mathcal{I}$ omits, which *is* mode 1. So all four modes reduce to threats against the identification or catalogue-completeness — the two claims with genuine empirical content — and none threatens an analytic consequence.

---

13. What the Paper Establishes

Established (conditionally), with the conditions named:

- Particle identity *can be reduced* to stable representation classes — the *sufficiency* direction (§5.1), conditional on fold-exhaustion, the Identification Premise (catalogue-completeness, §3), and the restricted Refinement–Transport Bridge (§4).
- **The central result is ontological (§9B):** particle species add no primitive beyond representation classes and invariant structure — the reduction from particle identity down to closure structure is strictly ontology-subtractive, removing a primitive at every rung and adding none. This is the paper's genuine contribution.
- **The invariant tuple carries the identity burden (§9A):** what looked like a heterogeneous list of particle properties (chirality, generation, holonomy, charge, confinement) is the structured content of one closure invariant. This is the substantive economy the reduction buys.
- **Species Carrier Uniqueness (§11):** given fold-exhaustion and catalogue-completeness, the representation class is the *unique adequate carrier* of particle identity. Two arguments converge: an *elimination* of named alternatives (§11.2), and a stronger *collapse dilemma* (§11.2A) showing any carrier whatsoever either fails observability (struck) or is observable (and collapses into an invariant, hence into the class). The dilemma **retires enumeration-completeness** as a load-bearing premise — the uniqueness no longer needs the list of alternatives to be complete — consolidating the carrier question onto catalogue-completeness alone. Not uniqueness simpliciter, and explicitly not the withdrawn "inexpressibility" claim.
- **The Operational Species Principle is physics-general, not VERSF (§11.4):** *given how particle physics already identifies species* — by measurable properties, with no surplus criterion — a species *is* an equivalence class of invariant observables. This owes nothing to the fold substrate; VERSF's distinctive and *separable* step is only the further identification of that class with a PFD class. A reader who rejects the fold ontology still owes the general principle. This two-layer split quarantines the controversial part.
- **Operational Species Criterion (§11.4):** under catalogue-completeness, two structures are the same species iff all admissible species-identifying observables agree, iff their invariant tuples are identical, iff they share a class. Both equivalences are *analytic given their premises*; the criterion inherits completeness's conditional, empirically-exposed standing rather than acquiring a stronger status from being conditioned — content lives in completeness being a falsifiable premise (§12, mode 1), not in the chain being more than an unpacking.

- **Physicists already use this criterion (§6A.1):** the field individuates species by agreement on measured quantum numbers and admits no surplus criterion — which *is* the Operational Species Principle's premise in practice. Near-unconditional, depending only on existing practice, not on the substrate premises. This is the paper's least controversial and arguably most persuasive claim.
- **Representation-Class Persistence Principle (§4.3):** under the identification, a species persists over exactly the interval it remains in its class — stated explicitly as a *conceptual consequence of the identification, not an empirical theorem* (both sides reduce to "the identity-defining invariants are preserved"). It fixes vocabulary; it discovers nothing about nature.
- The persistence interval is thereby *non-circular* and *empirical* (§4.2): class transitions ($\mu \to e$, $n \to p$) are observable invariant changes, so the interval boundary has physical content rather than being a definitional artifact.
- Stability is a *spectrum*, not a binary (§7): stable, metastable, and rapidly-decaying species differ only in interval length, all governed by the same persistence principle; the *form* of the spectrum is described (lifetimes are not).
- Representation classes provide operational indistinguishability (§6), at the Identification Premise's conditional standing — *the quotient structure only*, not exchange statistics.
- Species and objecthood are distinct questions (§4.4), related through the persistence principle of §4.3.
- Mass is downstream of identity (§10).

Reframed, not derived (the paper's reach as a gateway, §9C):

- exchange statistics → "what is the admissible exchange sector of a spinorial closure class?" (a structural question about class transport);
- generation count → "why does refinement persistence admit closure depths 1, 2, 3 and terminate there?" (a question about the range of one invariant, γ_D) — and, beyond reframing, the **candidate derivation** of §9.1: three dyadic sectors saturate the $K = 7$ register, a fourth overflows. All three premises conditional or open — P1 (budget) conditional on the hexagonal architecture's flavour-independence being *demonstrated*, P2 (dyadic law) open and narrowed by §9.2 to a branching-factor question, P3 (sum-rule) conditional. A candidate with a falsifier, not a result;
- flavour mixing → "what are the admissible transitions between species classes?" (a transition structure on the space of classes).

Each is stated as a *consequence* — if the identification holds, the problem becomes that question — importing no open result and asserting no derivation.

Not established (open or out of scope):

- the *necessity* direction — that only representation classes carry particle identity — is **argued** (§11), **not left open**, and after the collapse dilemma (§11.2A) it no longer rests on a separate enumeration-completeness premise: any carrier at all either fails observability or collapses into the class, so necessity is conditional on **catalogue-completeness alone** (the same premise the forward direction's deepest dependency rests

on). What remains genuinely open is therefore catalogue-completeness itself, not a distinct "have we named every alternative?" question;
- catalogue-completeness itself (§3), the failure mode being a found observable sensitive to a difference $\mathcal{I}$ omits;
- exchange statistics and the spin-statistics connection (§6), a **separate and larger open problem**;
- the gauge-sector interval-boundedness (§8), **(open)**: gauge identity is provisionally inherited through transport representation classes, with interval-boundedness the one open ingredient for the sector;
- individual particle lifetimes (§7), dynamical and downstream;
- exact Standard Model spectrum, masses, couplings, CKM/PMNS parameters, and the completeness of the particle dictionary (out of scope).

The residue, gathered: the reduction stands on three named conditionals — fold-exhaustion **(inherited, conditional)**, catalogue-completeness **(conditional)**, and the restricted bridge **(conditional)**. It no longer carries a separate "enumeration-completeness" open claim: the collapse dilemma (§11.2A) absorbed the discovered-carrier worry into catalogue-completeness, so the necessity direction is conditional on completeness rather than on having named every alternative. The ways the conditionals could fail are tabulated in §12, and the recast there is itself part of the paper's honesty: the empirically vulnerable claims are the *identification* and *catalogue-completeness*, not the §4.3 persistence principle, which is analytic given the identification and so cannot be falsified by any observation. None of the paper's content is an empirical prediction that could clash with experiment; the paper imports the prior ontology rather than modifying it, so it is observationally safe and internally exposed. This is stated as a limitation, not a virtue — see §13A and the genre note of §1: a framework that cannot be wrong because it does not reach far enough to be wrong has not yet made contact with measurement. The one place this paper *reaches* toward such contact is the generation-count candidate of §9.1 — which has the right shape (a generation-blind budget saturated by a growth law, with a fourth generation as falsifier) but is not closed, its growth-law and sum-rule premises remaining open. So the candidate is the programme's nearest approach to a falsifiable claim; it is not yet a closed one, and the roadmap of §13A.1 states exactly what would close it.

13A. What This Paper Does Not Reach

It is worth isolating, in one place and without hedging, the limitation that the genre note (§1) and the residue paragraph both gesture at, because it is the single most important thing a reader should carry away about the *status* of the result.

There is no point in this paper where the framework is forced to agree with a measurement, and nothing in it could have come out otherwise. Every result is either an identification this paper adopts (species = representation class, §5), an analytic consequence of that identification (§4.3), an inheritance from a prior VERSF paper (§2, §8, §10), or a reframing of an open problem (§9C). None is a quantity computed from the substrate that a measurement could confirm or refute. The paper's own list of disclaimers makes this exact point from the inside: no spectrum, no mass, no coupling, no lifetime, no mixing angle, no generation count, no exchange statistics. Collectively, those disclaimers say that the paper does not reach far enough into the world to be in danger from it.

This is the boundary between a *conceptual reduction* and a *physical theory*. A physical theory has at least one place where reality can push back — a number that comes out right or wrong (the electron mass, the Higgs mass, the fine-structure constant, three generations, the CKM hierarchy). This paper has none, by construction: it is about what a particle *is*, not what any particle's properties *are*. That is a legitimate and even valuable thing to be — ontological reductions have repeatedly preceded and enabled the physics that vindicated them — but it is not the same thing as a physical result, and the paper does not claim it is.

The honest scope of the contribution, then, is this. The paper offers a coherent, ontology-subtractive identification of particle species with representation classes of persistent closure structure, positions it against the existing literature on quantum identity (§6A), and shows how three open Standard Model problems become well-posed questions about class structure (§9C). What it does *not* offer is any *closed* forced contact with data. It does offer, in §9.1, a **candidate** generation-count derivation that has the right shape for such contact — an independently-fixed budget (K = 7, generation-blind) saturated by a growth law, with a fourth generation as falsifier — but that candidate is not closed: its growth-law premise (dyadic loading for arbitrary depth) and its sum-rule premise (shared occupancy) are open. So a reader assessing whether the VERSF programme is *physics* (as opposed to a self-consistent formal ontology) should look to whether §9.1's two open premises can be discharged generation-blind: if they can, the candidate becomes a theorem with a live falsifier, and *that* would be a physical result in a way the conceptual reduction is not. The candidate is the programme's nearest approach to data; it is not yet an arrival.

One consistency condition attaches to that claim, and is stated rather than assumed — and §9.1 is exactly where it bites. Calling the generation count the place the programme could reach a falsifier presupposes that $\gamma \leq 3$ is *not yet* a forced result. The §9.1 candidate confirms this is the current state: its growth-law premise (P2) is, at present, read off the three-term $\mathcal{D}$ = diag(1,2,4) — a bound that *presupposes three* — so "independently-fixed dyadic loading" remains aspirational and the target stays genuinely open. The budget premise (P1, K = 7) is by contrast *intended* to be independently fixed — and would, if its flavour-independence were demonstrated, give the candidate its non-trivial shape; but that independence is itself asserted rather than shown (§9.1, P1), so even the budget side is conditional, on a different and more tractable claim than P2. So the honest reading is the second branch of this condition: the existing generation structure rests, via P2, on input that presupposes three, and the candidate's contribution is to isolate *that* (and P1's independence) as the things to discharge generation-blind. If a companion paper later derives the dyadic law for arbitrary depth from refinement structure alone, *and* demonstrates the hexagonal architecture's flavour-independence, the candidate becomes a theorem and contact-with-data is achieved; until then it is a candidate with no fully-secured premise — P1 conditional on demonstrable priority, P2 open, P3 conditional. The paper takes no position beyond what §9.1 establishes, and states the condition so that the reader can see exactly which branch currently holds.

13A.1 What would convert this into physical support — a roadmap

Because the paper is candid that it has no measurement contact, it owes the reader the converse: a statement of what *would* supply it. The species reduction would gain genuine physical support — would cross from conceptual reduction toward physical theory — if it *enabled the derivation* of a quantity the Standard Model takes as input. Three targets are visible, in increasing distance from this paper:

1. **Generation count ($\gamma \leq 3$).** The nearest, and the only one with a built-in falsifier — and the one this paper takes furthest, in the §9.1 candidate and its §9.2–9.3 analysis. The candidate has the right shape: a budget ($K = 7$) saturated by a dyadic growth law, the two meeting at 7 by separate routes. What it lacks is on three fronts, each now stated as a definite task with a named success-and-refute condition (§9.1): P1's flavour-independence must be *demonstrated* — concretely, that the hexagon's *six* boundary channels are forced as the minimal complete cyclic transport cell, with no appeal to downstream multiplicity (refuted if "six" is ever justified by what it produces); the growth law must be *derived for arbitrary depth* (P2); and the shared-occupancy sum-rule must be argued — concretely, that generation sectors compete for a shared finite channel-pool rather than reusing channels (P3). §9.2 narrows P2 one step further: dyadic loading would follow if admissible refinement branches by exactly 2, and a two-sided squeeze (distinguishability-preservation forbids slower; closure-admissibility forbids faster) would force that branching factor — with the Fold's minimal-binary character the substrate-level reason to expect it, and a uniform-branching lemma (§9.2) the gap the squeeze still hides. So the concrete next targets are sharply posed: demonstrate P1's independence, prove the refinement-branching squeeze (closing P2, modulo the uniformity lemma), and argue shared occupancy (P3) — all generation-blind. Together they would convert the §9.1 candidate into a theorem predicting exactly three generations, refutable by a fourth.
2. **Exchange statistics.** Further out. The species reduction delivers indistinguishability (the configuration-space quotient, §6) but explicitly *not* the boson/fermion dichotomy, which needs the topology of the quotient and the spin-statistics connection. Physical support here would mean deriving *which* exchange representation a spinorial closure class realises — a result that would predict the observed statistics rather than accommodate them.
3. **Flavour hierarchy / mixing.** Furthest. The species reduction recasts CKM/PMNS structure as admissible transitions between representation classes (§9C), but the transition amplitudes — and their seed — are not derived. Physical support would mean computing a mixing element from closure geometry in a way that could have come out wrong against the measured value.

The roadmap is ordered deliberately: generation count first, because it is closest to the class structure this paper makes explicit and because it carries a clean falsifier; statistics and flavour further out, because each needs additional structure (spinorial topology; transition dynamics) the species reduction locates but does not supply. None is attempted here. But stating them gives the reader the answer to the fair question *"what would convince a physicist this is more than a self-consistent vocabulary?"* — and the honest answer is: not this paper, but a successor that turns one of these three reframed questions into a forced number. The species reduction's job is to have posed them as well-defined questions about class structure; converting any one into a prediction is the work that would earn the programme its physical standing.

Status: (this paper); a statement of limitation, included because the difference between a conceptual reduction and a physical theory is exactly the difference between what this paper does and does not do, and concealing that difference behind theorem-formatting would misrepresent the result. The roadmap (§13A.1) states what would close that difference, without claiming to close it.

---

14. Conclusion

Reduction Summary

For the reader who wants the whole result in one compressed statement:

Reduction Summary. Under (i) fold-exhaustion, (ii) catalogue-completeness, and (iii) the restricted Refinement–Transport Bridge, particle ontology reduces as

```
particle species
      ↓   (species identity ≡ class membership — §5, an identification)
representation class
      ↓   ([D] = { D′ : 𝓘(D′) = 𝓘(D) })
invariant tuple 𝓘
      ↓   (matter programme)
persistent closure structure
```

with species persistence described by $\mathcal{I}$-preservation (§4.3, a conceptual consequence of the identification — not a theorem), the *sufficiency* of representation classes argued (§5) and their *necessity* argued by the collapse dilemma (§11.2A — any carrier either fails observability or becomes a class, conditional on catalogue-completeness, no longer on an open enumeration), indistinguishability — but not exchange statistics — delivered (§6), and the whole chain shown to be ontology-subtractive (§9B). The first arrow, moreover, is owed to physics generally, not to VERSF: the Operational Species Principle (§11.4) makes "species = class of invariant observables" follow from how physics already identifies particles; VERSF adds only the bottom arrow, the identification of the class with persistent closure structure. The reduction adds no primitive at any rung; that, not any persistence biconditional, is the result.

That is the actual reduction: a chain from the observed species down to persistent closure structure, each arrow named and marked. The species-to-class arrow is owed to particle-physics practice (the Operational Species Principle, §11.4); the class-to-closure-structure arrow is

VERSF's own, separable, more-contestable step. Its worth is conceptual and ontological: it removes a category of primitive, not a discrepancy with data.

The Standard Model begins with particles. VERSF begins with persistent closure structure.

Under the assumptions stated explicitly here — fold-exhaustion, catalogue-completeness, and the restricted Refinement–Transport Bridge — stable particle species are carried by stable representation classes of Persistent Fold Defects, *sufficiently* (§5) and, by the collapse dilemma (§11.2A), *uniquely*: any rival carrier either fails observability or collapses into the class. That necessity rests on catalogue-completeness, not on an unprovable claim to have enumerated every alternative — the dilemma closes the list rather than requiring it complete. The paper keeps sufficiency and necessity distinct, and marks both conditional, rather than collapsing them into an unconditional theorem.

The contribution is therefore a *conditional reduction* of particle ontology, not a derivation of the spectrum. Particles cease to be primitive ingredients. They become persistent, identifiable, refinement-stable classes of closure structure whose observable properties — mass, coupling, flavour, transport behaviour — emerge through subsequent dynamics, and whose identity is fixed, prior to all of that, by the invariant tuple.

The paper's distinctive contribution, beyond inheriting the companion framework, is *ontological*, and it is best stated plainly: particle species are not primitive entities but representation-level classifications of persistent closure structure, and the reduction that establishes this adds no primitive at any rung (§9B). That is the genuine result. The persistence principle of §4.3 — that a species persists over the interval it remains in its class — is a conceptual consequence of the identification, useful for fixing vocabulary but analytic, not a discovery; an earlier draft of this paper overstated it as a proven biconditional and that overstatement is withdrawn (§4.3). The contribution is conceptual and ontological, not phenomenological: the paper buys an economy of primitives, not agreement with a measurement.

The full layering the programme exhibits is:

```
Fold
  ↓
Fact
  ↓
Record
  ↓
Persistent Fold Defect
  ↓
Representation Class          ← classification by invariant tuple 𝓘
  ↓   ⟸ THIS PAPER identifies this step: ⟹
Particle Species              ← species identity ≡ class membership (§5)
  ↓
Mass / Coupling / Flavour     ← downstream dynamics (§10)
```

This paper concerns the highlighted step — the *identification* of representation class with particle species — and nothing above or below it. The rungs beneath are inherited; the rungs above are dynamical and out of scope. The word is "identifies," not "derives": a representation class *is what a particle species is*, under the identification this paper adopts and argues the coherence and consequences of. No new physics is derived at the arrow; an ontological reduction is performed there.

A particle is not a starting ingredient.

A particle species is a representation-level classification of persistent closure structure — adding no primitive beyond the invariant tuple that fixes it, and individuated from its fellows by nothing, because class-identity is identity of every invariant.