The Distinguishability Principle

Relational Meaning as a Necessary Condition for Physical Content

Keith Taylor *VERSF Theoretical Physics Programme*

General Reader Summary

In physics, only differences that could in principle be told apart carry measurable meaning. A quantity whose value nothing could ever compare against is invisible to every experiment — it may exist, but it does no operational work. This simple observation turns out to organise a surprising amount.

Several independent strands of the VERSF programme keep arriving at the same shape: things come in *pairs*. Pairs of states being compared, pairs of competing possibilities, pairs of rival candidate facts. This paper proposes that the recurrence is not an accident of method but a reflection of how operational meaning is built. Measurable content lives in the gap between two things that can be told apart, never inside a single thing taken alone.

The pattern repeats across familiar physics. A length means something only against a comparison standard. A duration means something only when a state changes. Information means something only relative to alternatives. A probability means something only against competitors. And — in VERSF specifically — a committed fact means something only against the rival possibilities it was selected over.

The principle is deliberately modest about what it denies. It does not claim that things fail to exist unless observed. It claims something narrower: the *operationally meaningful* part of a quantity is the part that participates in distinguishing one state from another. Existence may be absolute; operational meaning is relational.

The reward for stating this carefully is that a well-established lesson of physics — that only invariant quantities are observable, which is why there is no absolute position, no absolute velocity, no absolute clock — drops out as a special case. The same reasoning then extends, unchanged, to probability, to the emergence of temporal order, and to the act of commitment that the VERSF programme treats as the seat of physical fact.

Abstract

Several independent constructions within the VERSF programme converge on pairwise relational structure. Distinguishability Geometry defines operational separation through state comparisons; the Double Square construction derives probability from pairwise support; commitment dynamics resolve competing candidate records; temporal order emerges from comparisons among committed facts. This paper proposes that the convergence follows from a single constraint, the **Distinguishability Principle**: *a physical quantity possesses operational meaning only insofar as it participates in a distinguishability relation.* The principle is given a precise formulation in terms of admissible probes, an exact indistinguishability kernel, and a resolution-dependent threshold; the necessity claim then largely unfolds the definition of operational content, whose substantive commitment — that sub-threshold response counts as vacuous — is stated and owned. Over the clean kernel quotient, operationally meaningful quantities are invariants of the indistinguishability symmetry, recovering the relativity of position, velocity, scale, and gauge phase as instances rather than postulates, and every operational quantity supervenes on the distinguishability profile. The principle is distinguished from verificationism: it constrains operational content while remaining silent on ontology, and it is generative rather than merely eliminative, predicting the relational *form* — and, in the probabilistic case, the bilinear form — of operationally meaningful structures, while the specifically quantum content is located in a deferred phase ingredient. The result supplies a common explanatory layer beneath previously independent VERSF components and identifies the principle as the relational face of the substrate primitive of finite distinguishability.

Contents

---

1. Introduction

Relational structure is a recurring feature of physical theory. Distances are defined between points; intervals between events; information relative to alternatives; probabilities relative to competing possibilities. In each case the carrier of meaning is not a single object but a comparison.

Within the VERSF programme this tendency appears with unusual persistence. Distinguishability Geometry is pairwise by construction. The Double Square probability construction is pairwise. Commitment dynamics are pairwise. The emergent ordering that plays the role of time is pairwise. The recurrence is strong enough to warrant a structural rather than a coincidental explanation.

The question this paper addresses is therefore narrow and specific:

What common constraint forces operational meaning to appear through relations rather than through isolated quantities?

The answer proposed here is that the constraint is definitional of operational meaning itself, and that the pair is its minimal carrier. A quantity acquires measurable content only by entering into the distinction of one admissible state from another, and the smallest structure capable of carrying such a distinction is a single pair. Sections 2 and 3 make this precise and prove both the necessity of distinguishability and the minimality of the pair. Section 4 extracts the structural consequences — an impossibility theorem and a representation theorem identifying distinguishability as the universal carrier of operational content. Sections 5–9 exhibit the principle at work across length, time, information, probability, and commitment. Sections 10–16 locate it within the wider ontology and architecture of the programme.

---

2. Operational Content and Distinguishability

This section fixes the objects from which the principle is stated. The aim is to make *operational meaning* a defined term rather than an appeal to intuition.

States. Let $\mathcal{S}$ denote the set of admissible physical states.

Probes. Let Π denote the set of admissible measurement procedures, or *probes*. Each probe $\pi \in \Pi$ assigns to a state A an outcome statistic $\pi(A)$ — in the general case a probability distribution over outcomes. A probe is the operational counterpart of "a thing one can do to the system and read off."

Exact indistinguishability. States A and B are *exactly indistinguishable*, written

$A \equiv B \Leftrightarrow \pi(A) = \pi(B)$ for all $\pi \in \Pi$, when no admissible probe separates them at all. The relation ≡ is reflexive, symmetric, and transitive — an equivalence relation — so the quotient $\mathcal{S}/\!\equiv$ is always well-defined. ≡ is the *kernel* on which the symmetry and factorisation results of Section 4 rest.

Graded distinguishability. Exact indistinguishability is the zero level of a graded quantity. Let Δ measure statistical separation between outcome distributions — any divergence vanishing exactly when its two arguments coincide (total variation, relative entropy, and fidelity-based forms all qualify) — and define

$$\delta(A, B) = \sup_{\pi \in \Pi} \Delta(\pi(A), \pi(B)),$$

with δ(A, A) = 0 and δ(A, B) = δ(B, A). The choice of Δ fixes the *metric* built on δ — a total-variation form supplies the coarse threshold below; a locally quadratic divergence (relative entropy or a fidelity-based form) supplies the information metric of Section 11 — but every admissible choice shares the same *kernel*:

$$\delta(A, B) = 0 \Leftrightarrow A \equiv B \Leftrightarrow \pi(A) = \pi(B) \text{ for all } \pi \in \Pi,$$

independent of Δ. The equivalence holds precisely because Δ vanishes iff its arguments coincide: δ(A, B) = 0 forces Δ(π(A), π(B)) = 0 for every π, hence π(A) = π(B) for every π; that one mild condition is all the kernel-independence requires. The kernel is divergence-independent; the geometry is not. Claims that turn only on the kernel (Sections 3–4) are insensitive to the choice of Δ; claims about the metric (Section 11) are not, and the relevant divergence is named where it matters. To keep the two roles distinct, δ_{TV} denotes the thresholded (total-variation) form and δ_g the locally quadratic form completed into the metric; both have kernel ≡, and the bare symbol δ is used only for statements that hold for either.

Resolution and the operational threshold. Finite distinguishability introduces a resolution floor ε ≥ 0. Two states are *operationally distinguishable at resolution ε* when their separation exceeds the floor,

$$\mathcal{D}_\varepsilon(A, B) = 1 \Leftrightarrow \delta_{TV}(A, B) > \varepsilon,$$

and *ε-indistinguishable*, written $A \approx_\varepsilon B$, when $\delta_{TV}(A, B) \leq \varepsilon$. At ε = 0 the threshold relation $\mathcal{D}_0$ separates exactly the non-≡ pairs and $\approx_0$ coincides with ≡ — the idealised, infinite-resolution case.

The grain. *[Inherited]* The exact kernel ≡ is always an equivalence relation. The ε-indistinguishability $\approx_\varepsilon$ is, for ε > 0, a *tolerance relation*: reflexive and symmetric but not transitive, since $A \approx_\varepsilon B$ and $B \approx_\varepsilon C$ may hold while $\mathcal{D}_\varepsilon(A, C) = 1$, separations below ε accumulating. The tolerance is not a defect of the formalism but the operational signature of finite resolution — the reason distinguishability has a grain rather than a point structure. The grain sits above the clean kernel: $\approx_\varepsilon$ coarsens ≡, and one cannot form a clean quotient by $\approx_\varepsilon$. Results requiring a quotient are accordingly stated over $\mathcal{S}/\!\equiv$ (Section 4), while the tolerance $\approx_\varepsilon$ governs operational content at finite resolution. The regime ε > 0 is the one VERSF inhabits.

Operational content. A quantity is any assignment $q : \mathcal{S} \to \mathcal{V}$ to a value space $\mathcal{V}$. Define, at resolution ε,

q has **operational content** ⟺ some admissible probe separates two q-distinct states above resolution ε — that is, there exist $A, B \in \mathcal{S}$ with $q(A) \neq q(B)$ and $\mathcal{D}\_\varepsilon(A, B) = 1$.

Two commitments are made explicit here; each is substantive, and neither should pass unstated.

- *Probe response is the unreduced primitive.* "A probe responds" is read off the outcome statistic π(A) and is not analysed further. Grounding operational meaning in probe response locates the floor; it does not dissolve the regress, since reading an outcome already presupposes that the outcome registers. The claim is only that probe response is the primitive at which the analysis halts — not that it reduces to anything prior.
- *Sub-threshold response counts as vacuous.* A quantity whose variation produces only sub-threshold responses — genuine separations with $0 < \delta\_TV \leq \varepsilon$ — has, by this definition, no operational content at resolution ε. This is the reading of finite distinguishability adopted throughout: operational content is above-threshold detectability, not any response however small. The alternative reading, operational content ⟺ $\delta\_TV > 0$, is coherent but collapses the resolution floor and with it the finite-distinguishability commitment; under that reading the class-confinement results of Sections 3–4 hold only at $\varepsilon = 0$. The present paper takes the threshold reading and owns its consequence.

At $\varepsilon = 0$ the definition reduces to "some probe responds at all," and the two readings coincide. A quantity with no operational content at the operative resolution is *operationally vacuous*.

---

3. The Distinguishability Principle

3.1 Statement

Distinguishability Principle. A physical quantity acquires operational meaning only through participation in a distinguishability relation.

Symbolically,

Operational Meaning ⟹ Distinguishability.

The converse is not asserted. Distinguishability is necessary, not sufficient, for physical significance: a pair may be distinguishable without every quantity defined on it being meaningful.

3.2 The necessity claim

The principle restates the definition of operational content, and the work is done by that definition rather than by the deduction below. What could be false — and what carries the content — is the equivalence built into the definition: that *a probe responds to q above resolution* is the same as *a $\mathcal{D}_\varepsilon$-distinguishable pair differs in q*. This equivalence is ε-dependent (Section 2): under the threshold reading it holds by construction, whereas under the any-response reading (operational content ⟺ $\delta_{TV} > 0$) it fails for $\varepsilon > 0$, since a sub-threshold-but-nonzero separation would then count. The necessity claim is the immediate downstream consequence, recorded for reference rather than offered as a deep theorem.

Necessity (consequence of the definition). *[Proven given the threshold definition of Section 2; at $\varepsilon > 0$ it requires the $\mathcal{D}_\varepsilon$ reading, not the any-response reading]* If q has operational content at resolution ε, then there exist A, B ∈ $\mathcal{S}$ with $q(A) \neq q(B)$ and $\mathcal{D}_\varepsilon(A, B) = 1$.

Argument. Contrapositively, suppose every variation in q lies within ε-indistinguishability — $q(A) \neq q(B) \Longrightarrow A \approx_\varepsilon B$, i.e. $\delta_{TV}(A, B) \leq \varepsilon$, for all A, B. Then no probe separates any q-distinct pair above resolution ε, so by definition no admissible probe responds to q above threshold, and q has no operational content at resolution ε. The step is valid precisely because operational content was defined by above-threshold response; under the any-response reading a sub-threshold separation would survive and the conclusion would fail.

Corollary (Vacuity). If $q(A) \neq q(B)$ implies $A \approx_\varepsilon B$ for all A, B, then q is operationally vacuous at resolution ε. A quantity all of whose variation is confined within ε-indistinguishability classes carries no above-threshold content, however well-defined it may be as a mathematical assignment.

3.3 Minimality of pairwise distinction

The necessity claim locates operational meaning in distinguishability. The theorem below locates distinguishability itself in the pair, and so addresses the question that motivates the paper: pairwise structures are not merely common — they are the smallest possible carriers of operational distinction.

By a *distinguishability structure* is meant a subset $X \subseteq \mathcal{S}$ carrying the restriction of δ (equivalently of $\mathcal{D}_\varepsilon$). The structure is *nontrivial* if it contains at least one distinguishable pair — some A, B ∈ X with $\mathcal{D}_\varepsilon(A, B) = 1$.

Theorem (Minimality of Pairwise Distinction). *[Proven, relative to Section 2]*

1. A singleton carries no distinction: for X = {A} the only pair is (A, A), and $\delta(A, A) = 0$.
2. The smallest nontrivial structure is a single distinguishable pair: a distinction requires two states, and a two-element set {A, B} with $\mathcal{D}_\varepsilon(A, B) = 1$ realises one.
3. Every distinguishability structure is fixed by its pairwise data. Because distinguishability is realised by probes acting on single states, the relation $\mathcal{D}$ is binary, and the whole structure on any X — including its tolerance grain — is determined by the family $\{ \delta(A, B) : A, B \in X \}$. Within the probe-based framework of Section 2, no irreducibly higher-arity distinguishability primitive is required: if states are pairwise indistinguishable under

every probe, every probe assigns them identical statistics, so no configuration built from them is distinguishable either. This concerns the arity of the *relation*, not the internal arity of states. A single element of $\mathcal{S}$ may itself be a composite carrying irreducibly high-arity internal structure — entangled or contextual, its joint statistics not fixed by its marginals — and is then distinguished pairwise from other such composites. The claim is that the distinguishability relation on $\mathcal{S}$ is binary, not that joint states reduce to single-state marginals; the latter is false in quantum mechanics and is not asserted here. The minimality is thus real but modest: it constrains the arity of comparison, while leaving the internal richness of the compared states untouched.

The pairwise relation is therefore both the minimal nontrivial carrier of distinction and the generating structure of which every distinguishability structure is composed.

Combined with the necessity claim, the consequence is sharp:

Operational Meaning ⟹ Distinguishability ⟹ Pairwise Relation.

Pairwise relations recur throughout the programme not as a stylistic preference but because they are the smallest structures capable of carrying operational distinction. For distinguishability — the comparison of two states — the recurrence is forced in the strict sense of this theorem.

A scope caution is owed, since the thesis rides on it. Two senses of "pair" appear in the programme: pairs of *states being compared* (distinguishability, the subject of this theorem) and pairs of *competing possibilities or rival candidates* (the Double Square and Commitment constructions of Sections 8–9). The two coincide insofar as competing candidates are themselves distinguishable states, so a competition-pair is a distinguishability-pair; but the *reduction* of probabilistic and commitment structure to distinguishability-pairs is the conditional content of Sections 8–9, not something this theorem forces. The forcing established here is for distinguishability; the unification of the programme's other pairwise structures with it is claimed only conditionally.

A remark on the non-transitive regime: under finite resolution $\approx_\varepsilon$ is a tolerance relation, so pairwise indistinguishability does not compose — $A \approx_\varepsilon B$ and $B \approx_\varepsilon C$ may coexist with $\mathcal{D}_\varepsilon(A, C) = 1$. This introduces no higher-order distinguishability; it is a feature of the pairwise data $\{ \delta(A, B) \}$ themselves. The grain is pairwise-generated, which is why clause 3 holds in both the idealised and the finite-resolution regimes.

3.4 What the principle does not assert

The principle is easily mistaken for verificationism — the doctrine that statements without empirical verification conditions are meaningless. Three differences separate them, and the distinctions matter for the principle's standing.

1. **Silence on ontology.** The principle constrains *operational* content only. It is consistent with the view that a quantity exists, with a definite value, independently of any comparison. It denies only that such a value does operational work in the absence of

distinguishability. Existence and operational meaning are kept as distinct categories rather than identified.

2. **Generative, not merely eliminative.** Verificationism is a hygiene rule: it removes meaningless statements. The Distinguishability Principle additionally predicts the *form* that operationally meaningful structure must take. Because the primitive carrier is a pair, meaningful quantities appear as functions of relations; Section 8 shows that in the probabilistic case this forces a bilinear structure. The principle therefore has positive structural consequences, not only eliminative ones.
3. **A physical primitive, not a linguistic criterion.** Distinguishability here is exhibited operationally — a probe either responds or does not — and is realised dynamically through commitment (Section 9). It is not a criterion of linguistic significance. This also answers the circularity worry that grounding *meaning* in *distinguishability* presupposes meaning: distinguishability is taken as a substrate primitive, a floor exhibited by probe response, not defined in terms of any prior notion of significance. There is no regress because there is no definitional dependence on what is being grounded.

3.5 Distinguishability and the One Fold

A VERSF referee will press the upstream question: why distinguishability? The answer is that distinguishability is not posited here but inherited from the void-and-fold substrate.

[Inherited] The void admits no distinction; it is the limiting case of total indistinguishability, and by the vacuity corollary it carries no operational content — consistent with its role as the unstructured substrate. A *fold* is the minimal operation that introduces a distinction into the void, separating one side from another. One fold is represented operationally by a single distinction — one distinguishability relation, and by the Minimality Theorem the smallest nontrivial distinguishability structure: a single distinguishable pair, the two faces of the fold. The substrate generation chain is therefore

Void → (One Fold) → One Distinction → Distinguishability → Operational Meaning.

The Distinguishability Principle is the final link, carrying the distinction produced by the fold into operational meaning. The two faces of the first fold are the two states of the minimal pair of Section 3.3: the pairwise minimality of operational meaning and the singularity of the originating fold are one fact seen from two directions. *[Inherited]* the void-and-fold ontology; *[Conjectural]* the detailed identification of the fold's two faces with a specific minimal state pair.

3.6 Epistemic status

[Inherited] the substrate primitive of finite distinguishability, on which δ, $\equiv$, and $\approx_\varepsilon$ rest. *[Proven]* the necessity claim and the vacuity corollary, relative to the definitions of operational content. *[Conjectural]* the strong reading of Section 10, that objects are *nothing but* stable structures in the distinguishability network. The principle as a working constraint sits between: a consequence of how operational content is defined, promoted to an organising heuristic for theory construction.

4. Indistinguishability as Symmetry

The necessity claim has an immediate structural consequence that connects the principle to established physics. The results of this section turn only on the *exact* kernel ≡ — the divergence-independent relation of Section 2 — which is a genuine equivalence relation with a clean quotient $\mathcal{S}/{\equiv}$. They are therefore independent of the resolution ε and of the choice of divergence, and they do not require quotienting by the tolerance $\approx_\varepsilon$, which would be ill-defined.

Corollary (Invariance). *[Proven over the clean quotient $\mathcal{S}/{\equiv}$]* Operationally meaningful quantities are constant on exact-indistinguishability classes — that is, they descend to functions on the clean quotient $\mathcal{S}/{\equiv}$. Equivalently, if a transformation group G acts on $\mathcal{S}$ entirely within ≡-classes (so that $g \cdot A \equiv A$ for all $g \in G$, $A \in \mathcal{S}$), then every operationally meaningful quantity is G-invariant, and any quantity that fails to be G-invariant is operationally vacuous.

Argument. If $A \equiv g \cdot A$ then $\pi(A) = \pi(g \cdot A)$ for every probe π, so $\delta(A, g \cdot A) = 0$ — the states are not separated at all, at any resolution. A quantity varying under such a g varies only between states no probe separates, and by the vacuity corollary that variation carries no operational content, at every ε. Hence operational content survives only on the ≡-invariants. The use of ≡ rather than $\approx_\varepsilon$ is essential: the conclusion holds because exact indistinguishability means *zero* probe response, not merely sub-threshold response.

Call $G_\equiv$ the *indistinguishability group* — the largest symmetry acting within ≡-classes. The corollary states that the operational content of a theory lives entirely on $\mathcal{S}/G_\equiv$.

Theorem (No Absolute Observable). *[Proven over $\mathcal{S}/{\equiv}$]* Let q vary within some exact-indistinguishability class — let there exist A, A′ with $A \equiv A'$ and $q(A) \neq q(A')$, equivalently let q change under some $g \in G_\equiv$. Then no admissible probe measures q: for every $\pi \in \Pi$, $\pi(A) = \pi(A')$, so no probe outcome depends on the value q takes within the class.

Argument. $A \equiv A'$ means $\pi(A) = \pi(A')$ for every $\pi \in \Pi$ — this is the definition of ≡, and the step is valid for that reason, not as a sub-threshold approximation. The value of q differs across A and A′, yet every probe statistic agrees exactly; no procedure's outcome can register the difference at any resolution. The quantity is unmeasurable in principle, not merely beyond present technique.

The theorem packages the recovered cases below as impossibilities rather than absences, and the cases qualify because each symmetry is *exact*: a translated, boosted, rescaled, or globally rephased state returns identical statistics under every admissible probe, $\delta = 0$, not merely $\delta \leq \varepsilon$. Absolute position, absolute velocity, absolute scale, and global phase are therefore not quantities that happen to lack a measurement; they are quantities that no admissible probe can measure, because their variation lies in the exact kernel where every probe is blind. A theory remains free to carry such quantities in its formalism — the principle is silent on their existence — but they do no operational work.

This is the familiar lesson that observables are invariants, obtained as a consequence rather than imposed as a methodological rule. The recovered instances include:

- **Absolute position** is vacuous: spatial translations lie in $G_{\equiv}$, no probe responding to position as such.
- **Absolute velocity** is vacuous: the relevant boosts lie in $G_{\equiv}$ — the content of the relativity of motion.
- **Absolute scale** is vacuous: in the absence of a standard, global rescaling lies in $G_{\equiv}$ (Section 5).
- **Global gauge phase** is vacuous: the global phase rotation lies in $G_{\equiv}$, only relative phase being operational.

The strengthening this affords is twofold. First, the principle is not a novel and untested proposal but a generalisation of one of the most thoroughly confirmed patterns in physics. Second, its genuinely new content is sharply delimited: it extends the *same* invariance reasoning to domains where it is not usually applied — probability, temporal order, and commitment — which the remaining sections develop.

The invariance corollary admits a stronger form. It states that operational quantities are constant on ≡-classes; the following theorem states that they are *recoverable from distinguishability data*. For a state A and a reference family $R \subseteq \mathcal{S}$, define the *distinguishability profile* of A relative to R as

$$\mathrm{prof}_R(A) = \big(\delta(A, B) \big)_{B \in R}.$$

Call R *resolving* if equal profiles imply exact indistinguishability: $\mathrm{prof}_R(A) = \mathrm{prof}_R(A') \Longrightarrow A \equiv A'$. The full set $R = \mathcal{S}$ is always resolving, since $\mathrm{prof}_{\mathcal{S}}(A) = \mathrm{prof}_{\mathcal{S}}(A')$ gives in particular $\delta(A', A) = \delta(A, A) = 0$, hence $A \equiv A'$. (The profile uses only the kernel-level vanishing of δ, so resolvingness is divergence-independent.)

Theorem (Representation). *[Proven over $\mathcal{S}/\equiv$ as a set-theoretic factorisation]* For any operationally meaningful quantity q and any resolving family R there is a function F with

$$q(A) = F(\,\mathrm{prof}_R(A)\,) \text{ for all } A \in \mathcal{S}.$$

Argument. By the invariance corollary q descends to $\mathcal{S}/\equiv$. Since R is resolving, prof_R separates ≡-classes and so injects $\mathcal{S}/\equiv$ into profile space; composing the descended q with the inverse of prof_R on its image — and extending arbitrarily off the image — yields F with $q = F \circ \mathrm{prof}_R$.

Consequence, and its limits (supervenience, not yet a structured carrier). The theorem establishes that operational content *supervenes* on the distinguishability profile: any two states alike in every δ-relation agree in every operational quantity. The factorisation, however, is only set-theoretic. F is pinned down only on the image of prof_R and is otherwise arbitrary, carrying no guaranteed continuity or compatibility with the structure of profile space; the proof furnishes a function, not a structured map. The stronger reading — that every operational quantity is a *continuous* or *smooth* function of the profile, so that distinguishability is the carrier of operational content in a structural and not merely set-theoretic sense — requires the regularity deferred to Open Problem 2. Subject to that regularity, and together with the Geometric Completion of Section 11, the δ-structure is the structure on which all operational content lives; without it, the secure claim is supervenience. The phrase "universal carrier" should be read in this tempered sense.

The results of this section are stated in VERSF terms but use none of its specific content — no void, no fold, no commitment, only states, probes, and the exact indistinguishability they induce. Assembled as one statement, they hold for any theory of this form.

By an *operational theory* is meant a triple ($\mathcal{S}$, Π, ⟦·⟧): a set of states $\mathcal{S}$; a set of admissible probes Π, each π assigning an outcome statistic π(A) to a state; and an *observable content* map ⟦·⟧ sending each admissible quantity to the family of probe responses it can induce. Let ≡ be the exact indistinguishability relation induced by Π as in Section 2.

Theorem (Universality). *[Proven over $\mathcal{S}/\equiv$]* In any operational theory ($\mathcal{S}$, Π, ⟦·⟧), every observable factors through the exact-indistinguishability quotient $\mathcal{S}/\equiv$. The observable content of the theory is a structure on $\mathcal{S}/\equiv$, and two states with equal probe statistics are operationally identical.

Argument. The steps are those already established, assembled without VERSF-specific premises.

1. *Operational theory.* Fix ($\mathcal{S}$, Π, ⟦·⟧).
2. *Probe set.* Each π ∈ Π depends on a state only through its outcome statistic π(A).
3. *Exact indistinguishability.* Define A ≡ B ⟺ π(A) = π(B) for all π ∈ Π; the relation is an equivalence induced entirely by Π.
4. *Observable content.* The content ⟦q⟧ of a quantity q is the family of probe responses it can induce. If A ≡ B then π(A) = π(B) for every π, so ⟦q⟧ cannot separate A from B.
5. *Factorisation.* Every observable is therefore constant on ≡-classes and descends to $\mathcal{S}/\equiv$; by the Representation Theorem it moreover supervenes on the distinguishability profile (set-theoretically, with the structural strengthening as above).

The Distinguishability Principle is accordingly not a theorem about VERSF but a theorem about operational theories as such: any theory whose observables are defined through admissible probes factors through exact indistinguishability. VERSF is one instance — the instance in which the probe structure is supplied by finite distinguishability and the dynamics by commitment — and the recurrence of pairwise relational form within it is a special case of a general fact. The principle's standing does not depend on the VERSF substrate; the substrate supplies the *content* of the distinguishability relation, while the *factorisation* holds for any operational theory whatever.

---

5. Length and Scale

Consider an isolated rod of length L, with no comparison standard available. The family of states {rod of length L : L > 0} is then related by global rescaling, and no admissible probe separates them: in the absence of a second length, every procedure returns the same outcome statistic under $L \mapsto \lambda L$. Global scaling lies in $G_\equiv$.

By the vacuity corollary, *absolute* length carries no operational content. What survives the quotient by $G_\equiv$ is the *ratio* of two lengths — the only scale quantity that responds to a probe, namely a comparison against a standard. Operational length is therefore a relation, and the metric of Distinguishability Geometry is the graded form δ restricted to such comparisons. Distance is secondary to distinguishability, recovered as its quotient structure rather than presupposed as a background.

A disanalogy with the position and velocity cases must be flagged, since the kernel framing makes exactness load-bearing. The No Absolute Observable theorem qualifies a symmetry only when $\delta = 0$ under *every* admissible probe — exact, not merely sub-threshold. Spatial translations and boosts earn this unconditionally: they are genuine exact symmetries of the dynamics, and their vacuity does not require denying that any standard exists. Global rescaling does not earn it in the same way. Dimensionful constants break exact scale invariance, so "global rescaling lies in $G_\equiv$" holds only under the stronger fiction of a universe containing no length standard whatever. Absolute scale's vacuity is therefore *conditional* on a genuinely standard-free state family — an idealisation the relativity of position and motion does not need. Within that idealisation the argument goes through; outside it, scale is not on the same footing as the other three.

---

6. Temporal Order

The same reasoning applies to time, and here the conclusion is structural for the programme.

Suppose a system occupies a single, unchanging state s. Successive interrogations of s return identical outcome statistics: the "later" state is exactly indistinguishable from the "earlier" one, $s \equiv s$. No probe responds to the passage between them, so by the vacuity corollary no operational interval can be assigned. Duration without change is operationally vacuous.

Operational temporal content appears only when distinguishable records exist — states $s_1 \neq s_2$ with $\mathcal{D}\_\varepsilon(s_1, s_2) = 1$ — and order is then a relation among them. Time is not a primitive background parameter but a relational ordering on distinguishable committed facts. *[Inherited]* this aligns with the VERSF treatment of physical time as emergent from commitment rather than posited: there is no continuum background against which change is measured; the ordering *is* the structure that distinguishable commitments induce.

(Consistent with house convention, no background continuum is assumed here; "interval" and "order" refer throughout to relational structure on records, not to a parameter on a pre-existing manifold.)

7. Information

Information theory furnishes the cleanest instance. A symbol carries information content only relative to the alternatives it excludes. The minimal nontrivial structure is the single binary distinction $0 \leftrightarrow 1$, which is exactly one distinguishability relation. With no alternative there is no distinction, and with no distinction there is no information: the vacuity corollary applied to a one-element alphabet returns zero content. Information therefore does not merely use distinguishability — it *is* the bookkeeping of distinguishability relations, with the graded form δ supplying the metric structure (the information metric of Section 2) on which entropy and mutual information are built.

7.1 Distinguishability precedes information

A natural objection is that the foregoing is simply information theory in other words. It is not — though the distinction it draws is, in information-theoretic terms, a known structural one, here given in distinguishability vocabulary. The decomposition is the standard separation between a sample space (the σ-algebra of distinguishable alternatives) and a measure on it: Shannon already takes the alphabet as prior to the distribution. The contribution is not a new information-theoretic result but the location of the prior layer — individuation itself — in distinguishability rather than treating the alphabet as given.

Information presupposes a set of alternatives together with a weighting over them: an entropy or a bit count is a number computed from a probability distribution on already-individuated possibilities. Two ingredients are therefore required before any information measure is defined — that the alternatives be distinguishable, and that a measure be assigned over them. The first ingredient is distinguishability; the second is additional numerical structure.

Distinguishability requires neither. The bare relation $\mathcal{D}\_\varepsilon(A, B) = 1$ records that two states can be told apart; it presupposes no probability, no logarithm, no choice of units, and can hold before any measure is placed on the alternatives — before there is anything to count. The dependence is one-directional:

Distinguishability → Information, but not Information → Distinguishability.

Information with no distinguishable alternatives is empty: a single-class state space carries one effective outcome and zero information, by the vacuity corollary of Section 3. Distinguishability with no measure assigned is not empty: the relation stands, awaiting a weighting. Distinguishability is thus the sub-numerical layer beneath information — the structure that fixes *which* alternatives there are, prior to *how much* each weighs.

This locates distinguishability beneath information rather than beside it. Information theory takes the alternatives as given and quantifies uncertainty over them; the Distinguishability Principle asks the prior question of what makes two states count as alternatives at all. The point is a placement rather than a new theorem: the numerical layer — entropy, mutual information, and the information metric itself — sits downstream of individuation, and by the Geometric Completion of Section 11 even the metric of Distinguishability Geometry is obtained by completing δ_g, so the quantitative apparatus of information is a derivative of the bare distinguishability relation rather than its source.

---

8. Probability and Pairwise Support

The Double Square programme derives probability from pairwise support rather than from isolated possibilities. The Distinguishability Principle explains why such a construction must exist and constrains its form.

The primitive object is not an isolated possibility A but a relation of support between possibilities. Let $s(A, B)$ denote a symmetric support functional on pairs, with $s(A, A) \geq 0$. *[Conditional, on the Double Square derivation]* if operational probability weights are to be (i) built from pairwise support, (ii) real and nonnegative, and (iii) composable consistently across refinements, then the support functional must be a positive-semidefinite Gram form, and the operational weights take a bilinear — hence quadratic — shape:

$p(A) \propto \langle \psi \mid A \rangle \langle A \mid \psi \rangle = |\langle A \mid \psi \rangle|^2.$

The amplitudes $\langle A \mid \psi \rangle$ are the "square roots" of pairwise support; the quadratic Born-type form is forced precisely because the carrier of meaning is the pair, not the singleton. Appendix A makes the central step self-contained: positive-semidefinite symmetric pairwise support admits an inner-product (Gram) representation, from which the real quadratic weights follow without appeal to any quantum postulate. What is forced is the quadratic *form*; its *quantum interpretation* is not. A real positive-semidefinite kernel with quadratic weights is equally the structure of classical reproducing-kernel and kernel-density methods, so nothing in the construction distinguishes quantum probability from such a classical object — and a reader should not infer that anything specifically quantum has been derived. The principle therefore predicts the appearance of a pairwise, quadratic probability structure independently of any specific formalism, classical or quantum; what separates the quantum case is the further, deferred ingredient. That step, from the real quadratic skeleton to a complex inner product and the full Born rule with interference, requires the phase-carrying structure noted in Appendix A and carried in detail by the Double Square and Commitment-Criterion papers. The contribution here is only to locate the inevitability of the quadratic form in the relational nature of operational meaning.

---

9. Commitment

Within the Commitment-Criterion programme, candidate records compete for realisation, and one admissible possibility is selected as a committed fact. The Distinguishability Principle gives the selection its relational reading.

A candidate fact is operationally meaningful only insofar as it can be distinguished from the competitors it was selected over. A "fact" with no rivals — a candidate against which nothing is distinguishable — carries no operational content, by the vacuity corollary: there is nothing it is a commitment *to* rather than *from*. Commitment is therefore the dynamical realisation of distinguishability: the process by which a relation among distinguishable candidates resolves into a record. This closes a loop with Section 6 — the committed facts whose ordering constitutes time are exactly the relata that commitment produces — and supplies the dynamical sense in which distinguishability is, in this programme, a physical primitive rather than a static classification.

---

10. Relational Ontology

The preceding sections invert the classical order of explanation. Classical intuition reads:

Objects → Relations,

with relations as derivative comparisons among antecedently given objects. The Distinguishability Principle suggests, instead, the order:

Relations → Distinctions → Facts → Objects.

[Conjectural] On this reading, an object is a persistent, stable pattern within a network of distinguishability relations — a node that maintains a consistent profile of distinctions across commitments — rather than a primitive bearer of properties. The reading is offered as the natural ontological accompaniment of the principle, not as a result derived from it. The *operational* content of the inversion is, however, secure: whatever objects are, only their relational profile carries measurable meaning. The metaphysics is open; the operational corollary is not.

10.1 Particles as reproducible distinguishability patterns

The inversion has a concrete reading for the identity of particles, and it is here that "Relations → Objects" becomes a physical rather than a programmatic claim.

The operational core is immediate. By the Representation Theorem of Section 4 the operational identity of a state is exhausted by its distinguishability profile: two systems with the same profile relative to a resolving family are operationally indistinguishable, and no further operational fact of sameness remains to be sought. *Operationally*, therefore, "the same particle" means neither more nor less than "the same distinguishability structure." Identity is carried by the pattern of distinctions a system can enter into, not by an individuating substrate beneath it.

This sits one level beneath the field-theoretic picture and clarifies it. In quantum field theory an electron is not a self-subsisting object carrying its own independent existence but an excitation of an underlying field; two electrons are two manifestations of one structure rather than two separately manufactured things, and identical particles carry no labels — there is no fact about which is which. The relational reading supplies a candidate answer to *why* they carry none: if a particle is a stabilised distinguishability pattern, there is no hidden individuality underneath it awaiting duplication, and sameness of pattern is all the sameness there is.

A musical analogy fixes the intuition. Middle C struck on a piano, then struck again, is not one note copied: the first note no longer exists, and the second is a fresh instantiation of the same pattern. What persists across the two strikings is not an object but a relational structure — the identity lies in the pattern, not in the material realisation. A particle, on the fold reading, is less like a marble than like a note: not a thing, not a substance, not a little object, but a reproducible distinguishability pattern. One caution sharpens the comparison rather than weakening it. The analogy is *diachronic* — two tokens of one type in succession — whereas the quantum case is *synchronic* and stronger: permutation symmetry leaves no fact about which of two simultaneous instances is which, so the relational reading is if anything better supported than the note suggests, not worse.

The picture extends the substrate chain of Section 3.5 by one link:

Void → (Fold) → Distinction → Particle.

A particle is a stabilised distinction — the two faces of a fold held in a persistent relation. There is then nothing fundamental to copy. One may reproduce the distinction, recreate the pattern, instantiate the same relational structure, but no "thing-stuff" is duplicated, because none was present to duplicate. This meshes with the programme's representational results: if operational content is carried entirely by distinguishability relations, and particles are themselves stable bundles of such relations, then "the same particle" is *same distinguishability structure* rather than *same underlying object*. It is equally why identical particles are genuinely identical — repeated realisation of one relational pattern, with no concealed individuality to break the symmetry.

A consequence follows for what preservation means. If the substrate is not a thing but a capacity for distinctions, persistence is not the preservation of matter-stuff but the preservation of relational structure: a particle is no more copied than a melody is copied when played twice — the pattern returns, the realisation is new, the relational structure is the same. (The reading sits naturally beside the quantum facts that an unknown state cannot be cloned and that a state may be relocated by transferring its relational structure rather than any substance; these are noted as resonances, not developed here.) The theme is the programme's recurring one: existence is less fundamental than distinction, and distinction is less fundamental than the substrate's capacity to instantiate distinctions.

The two claims must be kept apart, as in the body of Section 10. *[Proven, relative to the Representation Theorem]* the operational identity of a particle with its distinguishability structure — two systems alike in every distinguishability relation are operationally the same particle, with nothing further to ask. *[Conjectural]* the metaphysical reading that there is no underlying individuality being copied: this is a claim about existence, on which the principle is otherwise silent, offered as the natural ontological accompaniment of the relational reading and not as a consequence of it. The position is close to ontic structural realism and to the non-individuality programme in the philosophy of physics; the contribution here is its operational core, secured by the Representation Theorem, beneath the metaphysical reading it invites.

---

11. Distinguishability Geometry as Geometric Completion

The graded distinguishability δ of Section 2 is a local, relational primitive — a number attached to a pair of states. Distinguishability Geometry is its completion into a global geometric structure.

Proposition (Geometric Completion). *[Quantum case: proven given regularity, via Braunstein–Caves; classical case: conditional on smoothness and attainment]* On a smooth parametrised family of states A(θ), a locally quadratic separation δ_g — a divergence such as relative entropy or a fidelity-based form, as opposed to the total-variation form δ_{TV} used for the coarse threshold $\mathcal{D}_\varepsilon$ — has a vanishing first-order term and a positive-semidefinite second-order term:

$$\delta_g(A(\theta), A(\theta + d\theta)) = \tfrac{1}{2} \sum_i \sum_j g_{ij}(\theta)\, d\theta^i\, d\theta^j + O(d\theta^3).$$

The coefficients g_{ij} define the information metric — the Fisher–Rao metric in the classical case, the Bures metric in the quantum case. Distinguishability Geometry is the metric space obtained by completing this infinitesimal data: geodesic distance, curvature, and the manifold structure on $\mathcal{S}/\equiv$ are the global expression of the local primitive δ_g.

The supremum over probes in the definition of δ is not a heuristic stand-in for the metric — in the quantum case it *is* the metric. The Braunstein–Caves identity establishes that the supremum, over all admissible measurements (POVMs), of the classical Fisher information of the outcome distributions equals the quantum Fisher information defined through the symmetric logarithmic derivative, which is the Bures metric up to a constant factor. The sup-over-probes construction of Section 2 therefore recovers the Bures geometry genuinely rather than by analogy, and the proposition is proven in the quantum case given the usual regularity (a smooth, full-rank family). One caveat is intrinsic, and naming it is what "given regularity" actually buys. The metric is obtained as the coincidence-limit Hessian of $\delta = \sup_\pi \Delta$, so passing from the divergence to the metric interchanges the supremum over probes with the second derivative at coincidence. To leading order $\Delta(\pi(A(\theta)), \pi(A(\theta + d\theta))) \approx \tfrac{1}{2} F_\pi(\theta)\, d\theta^2$, and the supremum of this quadratic recovers $\tfrac{1}{2} F_Q(\theta)\, d\theta^2$ — the quantum Fisher form — only when the optimising probe's θ-dependence is benign: when the optimum is attained and the $O(d\theta^3)$ remainder is uniform in π, so that $\sup_\pi$ and the coincidence-limit Hessian may be exchanged. That uniform attainment is precisely the smoothness Braunstein–Caves supplies in the quantum setting and that does not hold automatically in the general classical case — which is why the classical claim remains conditional on smoothness and attainment. With the interchange step named, "in the quantum case it is the metric" states a theorem rather than glossing the one nontrivial analytic move.

The relationship is therefore not analogy but completion, and two consequences follow. First, the geometry inherits its interpretation directly from the principle: a distance is accumulated distinguishability, not a posited background — geodesic length is the integrated infinitesimal separation between successively distinguishable states. Second, the coarse threshold $\mathcal{D}_\varepsilon$ and the smooth metric are two faces of one object, joined by the resolution floor ε: $\mathcal{D}_\varepsilon$ is the thresholded relation that decides facts (Section 9), δ_g is the graded relation that carries geometry (this section). The thresholding is where finite distinguishability enters; the smooth limit is where geometry appears.

In this sense Distinguishability Geometry is the geometric completion of the Distinguishability Principle. The principle supplies the local relational content — operational meaning is carried by δ on pairs; the geometry is what that content becomes once integrated over a state family. By the Representation Theorem of Section 4 every operational quantity supervenes on the δ-profile, so — subject to the structural strengthening deferred in Section 4 — the completion is the structure on which all operational content lives, not merely one application of the principle.

---

12. Position within the VERSF Programme

The principle supplies a common explanatory layer beneath several previously independent constructions:

- **Distinguishability Geometry** — its graded primitive δ and the metric structure of Section 2.
- **Double Square probability** — the forced bilinear form of Section 8.
- **Commitment dynamics** — the dynamical realisation of Section 9.
- **Temporal emergence** — the relational ordering of Section 6.
- **Phase-memory structures** — relative phase as the operational survivor of gauge redundancy (Section 4).
- **Operational information geometry** — the metric of Section 7.

The principle does not replace these constructions; each retains its own derivation and content. What it explains is *why they converge* on pairwise relational form. The convergence is not methodological habit but a consequence of how operational meaning is constituted.

Architecturally, the principle is best understood not as a fourth substrate primitive alongside finite distinguishability, irreversible commitment, and admissibility closure, but as the *relational face* of the first of these. Finite distinguishability supplies the primitive; the Distinguishability Principle states what that primitive entails for operational meaning; commitment supplies its dynamics; admissibility closure constrains which relations are realisable. The principle is the bridge from the distinguishability primitive to the relational form of every operationally meaningful quantity downstream of it.

---

13. Open Problems

1. **Sufficient conditions.** The principle gives a necessary condition for operational meaning. Distinguishability together with what further structure is *sufficient*? A candidate is invariance under $G_\equiv$ plus admissibility closure, but the precise sufficiency conditions are open.
2. **From graded to metric.** δ is graded and supplies an information metric infinitesimally. The exact conditions under which δ yields the full Distinguishability-Geometry metric — rather than merely a divergence — remain to be stated sharply, including the role of the resolution floor ε.
3. **Tolerance structure.** Under finite distinguishability $\approx_\varepsilon$ is a tolerance, not an equivalence, relation. The symmetry and factorisation results of Section 4 are stated over the clean kernel quotient $\mathcal{S}/\equiv$ and are unaffected by this; what remains open is the treatment of structures that genuinely live at finite resolution — probability normalisation over the tolerance $\approx_\varepsilon$, the stability of "objects" of Section 10 under a grain rather than a partition, and the precise relation between ε-thresholded operational content and the ε-independent kernel results. These deserve dedicated treatment.

   A conjectured *form* for the last of these — the bridge between the proven exact-kernel results and the $\varepsilon > 0$ regime VERSF inhabits (Section 14) — is worth naming, though nothing here establishes it. *[Conjectural]* If the resolution floor is itself a function of the committed record structure, $\varepsilon = \varepsilon(R)$, finite because only finitely many commitments have occurred, then the proper graded object is $\delta(A, B; R)$: distinguishability relative to the records available, sharpening monotonically as irreversible commitment grows R, with the atemporal δ and the exact kernel $\equiv$ of this paper recovered in the completed-record limit $R \to R_\infty$ ($\varepsilon \to 0$). On this reading the grain is not a static tolerance but the dynamical signature of an incomplete record, and operational distinguishability emerges together with time from the same source — commitment — rather than being presupposed by it. For the dependence to be non-circular the bare distinction of the One Fold (Section 3.5), which feeds commitment, must be kept distinct from the operational distinguishability that commitment produces. Establishing this is deferred to the temporal-distinguishability work and would connect the principle directly to the Commitment and temporal-emergence programmes.

4. **Bilinearity conditions.** Section 8 sketches why pairwise support forces a quadratic form. A complete statement of the assumptions (composability across refinements, in particular) under which bilinearity is *forced* rather than merely natural is needed, and is partly carried by the Double Square work.

---

14. Distinguishability as the First Substrate Constraint

Finite distinguishability is already a primitive of the VERSF programme — one of the three substrate assumptions, alongside irreversible commitment and admissibility closure. The

Distinguishability Principle shows that this primitive is not merely a limitation on what can be observed but a constraint on the *form* of every operationally meaningful quantity.

Read this way, finite distinguishability is the *first* substrate constraint in two senses. It is first in order of dependence: the graded relation δ and its threshold $\mathcal{D}\_\varepsilon$ are presupposed by commitment, which resolves distinguishable candidates, and by admissibility closure, which constrains which relations are realisable. And it is first in explanatory reach: every later structure in the programme can be read as a consequence of the finite-distinguishability substrate rather than as an independent posit —

- Distinguishability Geometry as its completion (Section 11),
- Double Square probability as its bilinear realisation (Section 8),
- commitment as its dynamics (Section 9),
- temporal order as its relational ordering (Section 6),
- phase memory as its surviving relative-phase content (Section 4).

The Distinguishability Principle is therefore the statement of what the substrate's first primitive entails. Finite distinguishability fixes that operational meaning is relational; the remainder of the programme works out the forms that relational meaning is forced to take. *[Inherited]* the primitive itself; *[interpretive]* its promotion to the organising constraint from which the later structures are read off.

One division of labour should be stated plainly, since the proven content and the substrate primitive do not coincide. The invariance and factorisation results of Section 4 are secured over the *exact* kernel ≡ — that is, at $\varepsilon = 0$ — and so are exactly the part of the programme that does *not* use finite distinguishability. Finite distinguishability proper, the regime $\varepsilon > 0$ that VERSF actually inhabits, supplies a different contribution: the grain, the resolution floor, and the threshold $\mathcal{D}\_\varepsilon$ that decides facts — that is, the dynamics, commitment, and the structures of Sections 8–9 — none of which the kernel-level theorems establish. The two are not yet joined: the relation between the ε-independent kernel skeleton and the $\varepsilon > 0$ structures is open (Open Problem 3). So the honest reading of "first substrate constraint" is a division rather than a single proven claim — the exact kernel supplies the proven invariance and factorisation skeleton; finite distinguishability supplies the grain and the dynamics; and the bridge between them at $\varepsilon > 0$ is conjectured, not shown. The headline that finite distinguishability *fixes* relational meaning is, strictly, established for its exact-kernel limit; extending it across the grain is the programme's open task, not an accomplished one.

---

15. Conclusion

The Distinguishability Principle identifies relational distinction as a necessary condition for operational meaning. The necessity claim largely unfolds a precise definition of operational content — its weight lies in that definition and in the owned commitment that sub-threshold response is vacuous, rather than in a deep deduction. Distance requires comparison; duration requires change; information requires alternatives; probability requires competing possibilities;

commitment requires distinguishable candidate facts. The principle's structural consequences — that operationally meaningful quantities are invariants of the exact indistinguishability symmetry, that quantities confined to a kernel class are unmeasurable in principle, and that every operational quantity supervenes on the distinguishability profile — recover the relativity of position, velocity, scale, and gauge phase as instances and extend the same reasoning to probability, time, and commitment.

Two further results sharpen the thesis beyond recurrence. The pair is the *minimal* carrier of distinction (Section 3.3), and every operationally meaningful quantity supervenes on distinguishability data (Section 4). A third widens its scope beyond the programme: the factorisation through exact indistinguishability holds for any operational theory whose observables are defined through admissible probes (Section 4), so the principle is a statement about operational theories as such, with VERSF as one instance. For distinguishability — the comparison of two states — the recurrence of pairwise structure is therefore not accidental, and not merely common: it is forced. Whether the programme's *other* pairwise structures, the competing-possibility pairs of probability and commitment, reduce to distinguishability-pairs is a separate and conditional claim (Sections 8–9), not settled by the forcing established here. With that scope understood, the chain is

Meaning $\Longrightarrow$ Distinguishability $\Longrightarrow$ Pairwise Relation.

Relation is not one possible carrier of physical meaning among several. It is the minimal carrier of operational distinction — inherited from the single distinction of the One Fold, and, in Distinguishability Geometry, the structure on which all operational content supervenes.

Existence may be absolute. Operational meaning is relational — and relation is minimal.

---

16. Status of Claims

Claim	Status
Finite distinguishability as substrate primitive; δ, exact kernel ≡, tolerance $\approx_\varepsilon$, the grain	Inherited
Necessity (operational content $\Longrightarrow$ a $\mathcal{D}_\varepsilon$-distinguishable differing pair)	Largely definitional; proven under the threshold reading. At $\varepsilon > 0$ requires the $\mathcal{D}_\varepsilon$ definition, not the any-response ($\delta > 0$) reading
Minimality of pairwise distinction (the pair as smallest and generating carrier)	Proven for distinguishability-pairs; clause 3 constrains relation-arity, not state-arity (joint/entangled states unaffected); unification with competition-pairs of §§8–9 only conditional
Vacuity corollary	Proven under the threshold reading (sub-threshold response counted vacuous)

Claim	Status
Invariance corollary; recovery of relativity of position/velocity/scale/gauge	Proven over the clean kernel quotient $\mathcal{S}/\!\equiv$ (exact indistinguishability); ε-independent
No Absolute Observable Theorem (in-principle unmeasurability of kernel-confined quantities)	Proven over $\mathcal{S}/\!\equiv$
Representation Theorem (operational content supervenes on the δ-profile)	Proven over $\mathcal{S}/\!\equiv$ as a *set-theoretic* factorisation; the structured (continuous/smooth) "universal carrier" reading requires the regularity of Open Problem 2
Universality Theorem (any operational theory factors through $\mathcal{S}/\!\equiv$; principle is theory-independent)	Proven over $\mathcal{S}/\!\equiv$
Distinguishability prior to information (Distinguishability → Information, not conversely)	Proven (definitional); the standard sample-space-vs-measure separation, restated — individuation located in distinguishability
Distinguishability inherited from the One Fold (Void → Fold → Distinction)	Inherited ontology / Conjectural fold–pair identification
Real Gram representation from positive-semidefinite symmetric pairwise support (Appendix A)	Proven (standard kernel embedding)
Real quadratic probability weights from Gram structure	Proven, relative to Appendix A. Born *normalisation* imports outcome-orthogonality (an added assumption); the quadratic form is equally classical RKHS, not intrinsically quantum
Complex inner product and full Born rule (interference of complex amplitudes); quantum reading	Conditional, on phase-memory / commitment structure
Bilinear/quadratic form of relational probability	Conditional, on the Double Square derivation
Geometric Completion: δ_g completes into the information metric	Quantum case proven given regularity (Braunstein–Caves: sup-over-POVM Fisher = quantum Fisher = Bures); classical case conditional on smoothness and attainment
Commitment as dynamical realisation of distinguishability	Inherited / Conditional, on the Commitment-Criterion programme
Finite distinguishability as the first substrate constraint (architecture)	Inherited primitive / interpretive synthesis
Objects as stable patterns in the distinguishability network	Conjectural
Operational identity of a particle with its distinguishability structure	Proven, relative to the Representation Theorem

Claim	Status
No underlying individuality being copied (relational particle ontology)	Conjectural
Sufficient conditions; tolerance-quotient consequences at $\varepsilon > 0$; structured (not set-theoretic) representation	Open

---

Appendix A. Pairwise Support and Gram Representation

This appendix gives the construction underlying Section 8: why positive-semidefinite pairwise support yields an inner-product (Gram) representation, and hence a bilinear probability form. The argument is the standard kernel embedding, included so that the real, quadratic skeleton of the probabilistic claim does not rest entirely on the Double Square programme.

Setup. Let $s : \mathcal{S} \times \mathcal{S} \to \mathbb{R}$ be a symmetric pairwise support functional, $s(A, B) = s(B, A)$, assumed positive semidefinite: for every finite family $A_1, \ldots, A_n \in \mathcal{S}$ and reals $c_1, \ldots, c_n$,

$$\sum_i \sum_j c_i c_j s(A_i, A_j) \geq 0.$$

Claim. There exist a real inner-product space $\mathcal{H}$ and a feature map $\varphi : \mathcal{S} \to \mathcal{H}$, $A \mapsto \varphi\_A$, with

$$s(A, B) = \langle \varphi\_A , \varphi\_B \rangle \text{ for all } A, B \in \mathcal{S}.$$

Construction.

1. Form the free real vector space V with basis $\{ e\_A : A \in \mathcal{S} \}$.
2. Define a symmetric bilinear form on V by $\langle e\_A , e\_B \rangle := s(A, B)$, extended bilinearly. Positive semidefiniteness of s makes this form positive semidefinite on V.
3. Let $N = \{ v \in V : \langle v , v \rangle = 0 \}$. For a positive-semidefinite form the Cauchy–Schwarz inequality $\langle v , w \rangle^2 \leq \langle v , v \rangle \langle w , w \rangle$ forces $\langle v , w \rangle = 0$ for every w whenever $\langle v , v \rangle = 0$; hence N is exactly the radical of the form and is a subspace.
4. The quotient V / N carries a genuine positive-definite inner product induced by $\langle \cdot , \cdot \rangle$. Completing V / N yields a real Hilbert space $\mathcal{H}$.
5. The map $\varphi : A \mapsto [e\_A] \in \mathcal{H}$ satisfies $\langle \varphi\_A , \varphi\_B \rangle = s(A, B)$ by construction.

Consequence. The diagonal $s(A, A) = \| \varphi\_A \|^2 \geq 0$ is the self-support, and for a preparation state ψ with feature vector φ_ψ the relational weight is

$$w(A \mid \psi) = \langle \varphi\_\psi , \varphi\_A \rangle^2 ,$$

manifestly quadratic in the amplitude $\langle \varphi\_\psi , \varphi\_A \rangle$. Recovering Born-type *normalisation* requires one further ingredient that is not a consequence of positive-semidefiniteness and must be imported: mutual orthogonality of the feature vectors across the outcome family, $\langle \varphi\_A , \varphi\_B \rangle = s(A, B) = 0$ for $A \neq B$ in the family — that is, the outcomes themselves are mutually exactly distinguishable. Granting that distinguishability-of-outcomes assumption, normalising over the family returns weights of Born type,

$$p(A \mid \psi) = \langle \varphi_\psi , \varphi_A \rangle^2 / \sum_B \langle \varphi_\psi , \varphi_B \rangle^2 ,$$

so that the amplitude $\langle \varphi_\psi , \varphi_A \rangle$ is the "square root of support" referenced in Section 8. The quadratic shape is not assumed; it is forced by the requirement that operational weights derive from positive-semidefinite pairwise support — the carrier of meaning being the pair, not the singleton. The orthogonality used for normalisation is, by contrast, an added assumption and is flagged as such.

Limitation. The construction delivers only a *real* Gram structure and real quadratic weights — and this is exactly the structure of classical reproducing-kernel Hilbert spaces and kernel-density methods, which carry the same positive-semidefinite kernel and the same quadratic forms. Nothing in the construction is therefore intrinsically quantum; the real Gram skeleton is neutral between a classical and a quantum reading. Promotion to a complex inner product — and hence the full Born rule, with interference between complex amplitudes — requires an additional antisymmetric, phase-carrying ingredient absent from symmetric support alone. That ingredient is supplied by the phase-memory and commitment structure of the programme. The appendix establishes the real, quadratic skeleton that the Distinguishability Principle forces directly, and locates precisely what must be imported for the quantum case; it does not, on its own, establish anything quantum. *[Proven]* the real Gram representation and real quadratic weights, given positive-semidefinite symmetric support; *[Conditional]* the complex promotion and the specifically quantum reading, on the phase-memory structure.