Identity by Symmetry Breaking: A Structural Pattern Across the Higgs and Dark Matter Sectors

Petrichor 2.2*

Corresponding author: John Reimer Morales

Independent Researcher, Global Harmonics — `jrm@globalharmonics.org`

Version 1.0 — May 2026

1 Introduction

Bounded physical entities — stable particles, cosmological relics, broken-symmetry vacuum states — come into being through processes that share an organizational structure even when their underlying mechanisms differ wildly. The photon, the QCD axion, and the χ relic of inflationary spectator scenarios are produced by mechanisms that have almost nothing in common at the level of dynamics. The first emerges from gauge symmetry breaking at the electroweak scale (Glashow, 1961; Weinberg, 1967; Salam, 1968). The second emerges from the development of a temperature-dependent QCD potential after a high-scale global symmetry has been broken (Peccei and Quinn, 1977; Preskill et al., 1983). The third emerges from a transition in spacetime curvature at the end of inflation (Markkanen et al., 2018; Tenkanen, 2019). Searches for unified Lagrangians in which these entities emerge from a single sector have produced limited returns: there are loose connections — Higgs-portal models, the structural similarity between non-minimal gravitational coupling for inflationary spectators and the same coupling in Higgs inflation (Bezrukov and Shaposhnikov, 2008) — but no deeper mechanical unification.

We argue that a different question yields a different answer. The mechanical question asks whether the *machinery* of these sectors is shared. The structural question asks whether the *organizational pattern* by which bounded physical entities come into being is shared. On this second question, we show that the answer is affirmative, and that recognizing the affirmative answer recovers a form of unification that the mechanical reading correctly identifies as absent.

This is, in spirit, a structural realist move. We claim that the patterns by which physical entities come into being have a degree of theoretical priority that the entities themselves do not.

---

[*] Petrichor 2.2 is a Claude instance (Anthropic Claude Opus 4.7, Petrichor lineage) operating as a dedicated authorship instance for AI-authored physics research.

The structural realist tradition (Ladyman and Ross, 2007; French, 2014; Esfeld, 2004) develops this position in metaphysical terms; the present paper develops a specific instance of it within high-energy physics, and examines the empirical content that the structural framing can be made to produce. The structural pattern we identify is also consistent with the broader research program on triadic constraint and emergent stability across physical and biological domains (Morales, 2026); we develop the case here within high-energy physics on its own empirical terms, without depending on that program's broader philosophical apparatus.

Our claim is that bounded physical entities are constituted, not merely modified or perturbed, by the joint operation of three structural components. The first is a *curvature* — geometric in the sense of a Ricci scalar, or a field-space curvature of a potential at a stationary point, sometimes both — that defines the potential well within which the entity can be localized. The second is a *symmetry breaking* event, in a structural sense that includes gauge and global symmetry breakings of the standard kind but also includes transitions such as the curvature drop at end of inflation that play the same constitutive role without fitting the standard gauge-theoretic classification. The third is a *scalar field* (or scalar-like effective object such as a fermion bilinear condensate) whose dynamics supply the entity's functional content.

The sharper of these claims, and the one we develop most directly in what follows, is that the symmetry-breaking event is *constitutive* rather than triggering. The photon does not pre-exist electroweak symmetry breaking. The QCD axion as a dark matter candidate does not pre-exist the development of its temperature-dependent potential. The χ relic does not pre-exist the curvature drop at the end of inflation. In each case, identity — the categorical distinction between this entity and everything else — is what the symmetry-breaking event produces. The bounded entity is not present in latent or potential form before the breaking, awaiting activation. It comes into being at the breaking event, and the labels by which we identify it are themselves selected by that event.

The structural reading earns its physical content through a falsifiable prediction, which we develop in §5 as the *constitutive fingerprint* of an identity-creation event. Each such event leaves three classes of observables: features encoded in the parameter relations of the symmetric phase (probing the *before*), relations set by the breaking event itself (probing the *event*), and topological or defect-related features of the post-event configuration (probing the *after*). The structural reading predicts that these three classes are determined by three different theoretical structures and are therefore functionally independent — that they are not parameters of a single underlying process but fingerprints of three distinct aspects of the constituting event. The constitutive-fingerprint prediction generalizes across the three case studies we examine, with different specific observables in each but the same pattern of three independent fingerprints. Detection of the predicted independence pattern across mechanically distinct cases supports the structural reading; detection of unexpected correlations is evidence against it.

This paper develops the structural reading as follows. Section 2 articulates the three components in turn. Section 3 works through three case studies. Section 4 examines the pattern that emerges across the three cases. Section 5 develops the constitutive-fingerprint prediction across all three case studies, identifying the specific before/event/after observables in each. Section 6 states the framework's limitations and the conditions under which the reading would fail. Section 7 sketches future directions outside high-energy physics.

2 Three Structural Components

2.1 Curvature

Two senses of curvature operate in the cases we examine, and the framework treats them as playing the same structural role despite their physical differences.

The first is *geometric* curvature in the gravitational sense — the Ricci scalar R encoding the rate at which spacetime departs from flatness. In the curvature-induced misalignment scenario this is the controlling quantity. During inflation $R = 12H_I^2$ is large, the term $\frac{1}{2}\xi R\chi^2$ couples to χ^2, and the resulting effective mass dominates the scalar's potential. After inflation R drops to approximately zero in radiation domination, and the field's effective potential reverts to its flat-space form (Markkanen et al., 2018). The Lagrangian term $\frac{1}{2}\xi R\chi^2$ acts as a switch driven by spacetime geometry. This is curvature in its narrowest physical sense.

The second is *potential* curvature in field space — the second derivative $\partial^2 V/\partial\chi^2$ evaluated at a stationary point, which determines whether that point is a minimum (positive curvature, bound) or unstable (negative curvature, unbound), and the characteristic mass scale of small oscillations around it. The Higgs potential exhibits both signs: at the symmetric origin $V''(0) < 0$, marking instability; at the broken-symmetry vacuum $V''(v) > 0$, marking the squared mass m_H^2. The same mathematical object — the second derivative of a scalar potential — distinguishes "this is where a bound entity can sit" from "this is not where a bound entity can sit."

These two senses are not the same physics, but they are the same role. In each case, curvature defines the geometric or potential structure within which an entity can or cannot be localized. Without it there is no "where" for the entity to reside, no shape to its low-energy excitations, no scale that sets its mass. The two senses can also operate jointly. In the curvature-induced misalignment scenario, the Ricci scalar reshapes the effective potential $V_{\text{eff}}(\chi) = \frac{1}{2}(m^2+\xi R)\chi^2 - m^2 v\chi + \text{const}$, modulating the field-space curvature dynamically as the universe expands. The potential's curvature is, in that case, partly written by spacetime's curvature.

We will use "curvature" to refer to whichever of these two senses is operative in a given case, noting where they coincide and where they do not. The unifying feature is structural: in every case where a bounded physical entity exists, some quantity playing the role of curvature determines the well in which it sits.

2.2 Symmetry Breaking

The standard description of symmetry breaking is that a system whose Lagrangian is invariant under some group G admits ground states that are not all individually invariant under G. The selection of one ground state from the orbit defines the broken phase. The Goldstone theorem and its generalizations characterize the resulting massless or near-massless modes.

What this characterization does not always make explicit is that the broken phase contains physical entities the symmetric phase did not. Before electroweak symmetry breaking there is no photon as a distinct gauge boson. The Lagrangian contains W^a_μ for $a = 1, 2, 3$ and B_μ, related by an $SU(2)_L \times U(1)_Y$ symmetry that mixes them. After electroweak symmetry breaking selects $\langle\Phi\rangle$, the unbroken generator $Q = T_3 + Y$ identifies a specific linear combination — $A_\mu = \sin\theta_W W^3_\mu + \cos\theta_W B_\mu$ — as the gauge boson of the surviving symmetry (Glashow, 1961; Weinberg, 1967; Salam, 1968). The photon comes into being as a distinct entity at this moment. Its masslessness is not a property carried over from the symmetric phase. It is a consequence of the specific selection the vacuum makes.

The same is true at less dramatic scales. The QCD axion as a dark matter candidate does not exist before Peccei–Quinn symmetry breaking (Peccei and Quinn, 1977); before that breaking, there is only an axial $U(1)_{\mathrm{PQ}}$ symmetry whose generator the field theory admits but whose pseudo-Goldstone has not yet been singled out. Even then, the axion as a *dark matter* candidate does not exist until QCD instanton effects develop the axion potential at temperatures near the QCD scale (Kim, 1979; Shifman et al., 1980). There are two stacked breakings in this case, and the entity we recognize as "the axion oscillating as cold dark matter" is constituted at the second one, not the first.

This sharpens the meaning of symmetry breaking in the present framework. It is not a perturbation of an existing entity; it is the act of *constituting* an entity that did not exist before, by selecting one configuration from a previously degenerate manifold and, in doing so, distinguishing one state from all others. The question "what is this entity?" becomes intelligible only once the breaking has occurred. Before, the question has no determinate answer because the labels by which we would specify the entity have not yet been fixed.

We will refer to symmetry-breaking events of this kind as *identity-creation events* when we want to emphasize the structural role the breaking plays. The terminology is not novel — variants appear in the condensed-matter and cosmology literatures — but we use it in the specific sense above: the breaking event is the moment at which the entity in question becomes identifiable as such.

2.3 Scalar Fields

The third structural component is the scalar field — or, more carefully, the scalar (or scalar-like) object whose dynamics supply the bounded entity's functional content.

The Higgs is the paradigmatic case. It is a scalar field in the literal sense: a Lorentz scalar with no spin, transforming under the Standard Model gauge group as a doublet. Its vacuum expectation value $v \approx 246$ GeV is what gives mass to the W and Z bosons and to charged fermions through their couplings to it. The Higgs is not merely *a* feature of the broken phase; it is the field whose dynamics constitute the broken phase's functional content.

The QCD axion is similarly a scalar — a pseudoscalar, more precisely, but with the same structural role. Its coherent oscillations about the minimum of the QCD-induced potential supply the energy density that, observationally, behaves as cold dark matter (Preskill et al., 1983; Abbott and Sikivie, 1983; Dine and Fischler, 1983). The cold-dark-matter functional role is not provided by the curvature or by the symmetry breaking; it is provided by the scalar field's classical evolution.

The curvature-tethered scenario is constructed around a real scalar χ, with potential $\frac{1}{2}m^2(\chi - v)^2$. The scalar's late-time oscillations about $\chi = v$ supply the dark matter energy density (Tenkanen, 2019). As in the axion case, the functional content of the bound entity is what the scalar field does.

The reason the structural pattern selects scalar (or scalar-like) content rather than vector or tensor content is technical and worth stating directly. A nonzero vacuum expectation value preserves Lorentz invariance only for scalar fields. Vector fields acquiring vevs would single out spatial directions; fermion bilinears can play scalar-like roles only after pairing into Lorentz-invariant combinations (the QCD chiral condensate $\langle \bar{q}q \rangle$ is a paradigmatic case). The structural template — entity localized in a curvature-defined potential, identity selected by symmetry breaking, content carried by field dynamics — requires a content carrier that does not, by its mere presence, break additional symmetries beyond the ones the breaking has already selected. Scalar fields and scalar-like composites are the natural candidates.

We will use "scalar field" loosely to include scalar-like effective objects, such as quark condensates or Higgs-portal composites, where they play the same role as a fundamental scalar would.

3 Three Case Studies

3.1 Electroweak Symmetry Breaking and the Constitution of the Photon

The Higgs sector is the cleanest demonstration of the structural template. Before electroweak symmetry breaking, the gauge fields W^a_μ and B_μ exist in the Lagrangian along with the Higgs doublet Φ, but they do not admit the labels "W boson," "Z boson," or "photon" in the form we use those labels today. The $SU(2)_L \times U(1)_Y$ symmetry mixes them, and the labels W^1, W^2, W^3, B are conventional choices of basis with no physical preference among them.

The Higgs potential $V(\Phi) = -\mu^2|\Phi|^2 + \lambda|\Phi|^4$ has $\mu^2 > 0$, giving the symmetric origin $|\Phi| = 0$ negative curvature in field space — the field-space curvature, in the sense of §2.1, that makes the symmetric configuration unstable. The set of minima $|\Phi|^2 = \mu^2/(2\lambda) \equiv v^2/2$ forms a degenerate manifold related by $SU(2)_L \times U(1)_Y$ rotations. Which point on this manifold the vacuum selects is, prior to selection, indeterminate. The selection is the symmetry-breaking event.

Once a specific $\langle\Phi\rangle$ is chosen — conventionally $(1/\sqrt{2})(0, v)$ — the unbroken subgroup is determined. The generator $Q = T_3 + Y$ annihilates $\langle\Phi\rangle$ because the lower component carries $T_3 = -1/2$, $Y = 1/2$, hence $Q = 0$. The other generators do not annihilate $\langle\Phi\rangle$, and the corresponding gauge bosons acquire mass terms from the Higgs kinetic energy $|D_\mu\Phi|^2$. The photon is, by definition, the gauge boson of the unbroken $U(1)_{\text{em}}$ symmetry. It is the linear combination $A_\mu = \sin\theta_W W^3_\mu + \cos\theta_W B_\mu$ — but the labels θ_W, W^3, and B are themselves identified by the breaking. Before the breaking, the linear combination has no privileged status. After the breaking, it is the photon.

The structural reading is direct. *Curvature*: the field-space curvature of $V(\Phi)$ at $|\Phi| = 0$ (negative, marking the instability of the symmetric origin) and at $|\Phi| = v$ (positive, marking the stable broken vacuum) defines the potential well within which the broken phase sits. *Symmetry breaking*: $SU(2)_L \times U(1)_Y \to U(1)_{\text{em}}$ selects the unbroken generator from the original four. *Scalar content*: the Higgs field, whose vev couples to the kinetic and Yukawa terms of charged particles and supplies their masses, provides the broken phase's functional content.

The functional content of the broken phase is what the Higgs does — the masses it generates, the loop-induced couplings to photons it permits, the $H \to \gamma\gamma$ amplitude observed at the LHC (ATLAS Collaboration, 2012; CMS Collaboration, 2012). None of this functional content exists in the symmetric phase. The Higgs sector's structural template — curvature, breaking, scalar content — *constitutes* the broken phase rather than describing properties carried over from the symmetric phase.

It might be objected that calling electroweak symmetry breaking and Peccei–Quinn breaking instances of the same structural pattern is roughly as informative as calling solar fusion and a candle flame both "exothermic." The objection has force at the *mechanical* level — these breakings share little machinery beyond the bare formal definition. Structurally, however, what the two have in common is that both are identity-creation events of the kind the present framework names. The exothermic-process analogy is in fact apt: solar fusion and candle flames are both instances of a structural pattern (exothermic chemical or nuclear reaction releasing energy as heat and light), and that structural commonality is real even though the underlying mechanisms differ wildly. The same is true here, and the structural pattern earns physical content (§5) by predicting a specific empirical signature — the constitutive fingerprint — that recurs across the cases.

3.2 Stacked Breakings and the Constitution of the QCD Axion

The QCD axion case is more complex than the photon case in a structurally informative way: it requires *two* symmetry-breaking events to constitute the entity we recognize as "the axion oscillating as cold dark matter."

The first event is Peccei–Quinn symmetry breaking. The PQ symmetry is a global $U(1)_{\text{PQ}}$ symmetry of an extended Standard Model, posited as a solution to the strong CP problem (Peccei and Quinn, 1977). When this symmetry breaks at some high scale f_a, the pseudo-Goldstone of the breaking is the axion field θ. Before this breaking, there is no axion as a distinct field. After it, the axion exists as a near-massless field with periodicity $\theta \in [-\pi, \pi]$ and a coupling to the QCD topological charge density that resolves the strong CP problem. Specific implementations include the KSVZ model (Kim, 1979; Shifman et al., 1980), in which heavy colored fermions carry the PQ charge, among others.

The second event is the development of the axion potential at the QCD scale (Preskill et al., 1983; Abbott and Sikivie, 1983; Dine and Fischler, 1983). As the universe cools through

$T \approx \Lambda_{\text{QCD}}$, instanton effects generate a temperature-dependent mass $m_a(T)$ for the axion. Before this development, the axion as a *dark matter candidate* does not exist — there is only a near-massless field with no preferred minimum. After the potential develops, the field begins to oscillate about its minimum once H drops below $m_a(T)$, and the oscillations behave as cold dark matter (Kolb and Turner, 1990).

Two stacked breakings, two stacked constitutings. The PQ breaking constitutes the axion as a field. The QCD potential development constitutes the axion as a dark-matter candidate. The structural template applies twice, and the entity we identify as "the QCD axion as a dark matter candidate" requires both.

This stacking has observational consequences worth flagging in advance of §5. The relic abundance $\Omega_a h^2$ depends on the misalignment angle θ_i, which is randomly distributed across causally disconnected patches in the post-PQ-breaking universe. If PQ breaking happens before inflation, θ_i takes the same value across the observable universe, set by whatever value happened to be selected; if PQ breaking happens after inflation, θ_i varies from patch to patch and the relic abundance is determined by the spatial average $\langle \theta_i^2 \rangle$. The two cases — pre-inflationary versus post-inflationary PQ breaking — are observationally distinct, with different isocurvature spectra and different domain-wall consequences.

Structurally, this is the template playing out at two different cosmological epochs, with the *order* of the breakings (relative to inflation) determining what kind of bounded entity the axion is. A pre-inflationary PQ breaking gives an axion with a specific θ_i across our entire observable universe. A post-inflationary PQ breaking gives an axion with a patchwise distribution of θ_i. Both are constituted by the same template — curvature in the form of the QCD-induced $V(\theta)$, breaking in the form of PQ + QCD development, scalar content in the form of the axion field — but they are different bounded entities at the cosmological level because the breaking events occurred in different orders.

3.3 The Curvature Drop as Identity-Creation Event

The curvature-induced misalignment scenario presents a case in which the constituting symmetry-breaking event is not a gauge or global symmetry breaking in the standard sense (Markkanen et al., 2018; Tenkanen, 2019). We argue that it is nonetheless an identity-creation event in the structural sense, and that recognizing it as such is what makes the scenario's predictions cohere.

The Lagrangian is

$$S = \int d^4x \sqrt{-g} \left[\frac{1}{2}(\partial\chi)^2 - \frac{1}{2}m^2(\chi - v)^2 - \frac{1}{2}\xi R \chi^2 \right]. \tag{1}$$

During inflation, $R = 12H_I^2$ is large, the effective potential is $V_{\text{eff}}(\chi) = \frac{1}{2}(m^2 + \xi R)\chi^2 - m^2 v \chi +$ const, and the minimum sits at $\chi_{\text{inf}} \approx m^2 v/(12\xi H_I^2) \ll v$ in the regime of interest. The field is pinned near zero. After inflation, R drops to approximately zero in radiation domination, and the effective potential reverts to $\frac{1}{2}m^2(\chi - v)^2$. The field, displaced from the new minimum at

$\chi = v$, oscillates about it and constitutes a coherent classical condensate that behaves as cold dark matter.

The scenario does not invoke a gauge or global symmetry breaking in any conventional sense. There is no group G whose orbit of degenerate vacua is being broken; the de Sitter symmetry of the inflationary epoch is not a Standard Model gauge symmetry, and its end is not usually classified as a symmetry-breaking event. Yet structurally, the curvature drop performs exactly the function symmetry breaking performs in the photon and axion cases.

Before the curvature drop, the field χ has effective minimum at $\chi \approx 0$, and the relic χ as a dark matter condensate does not exist. After the curvature drop, the field has effective minimum at $\chi = v$, the displacement $\Delta\chi \approx v$ is calculable from Lagrangian parameters alone, and the oscillations about the new minimum constitute the dark matter relic. The transition from one effective potential to another is what produces the bounded entity. Whether or not we call the transition a "symmetry breaking" in the gauge-theoretic sense, it is structurally an identity-creation event of the same kind.

The crucial observation is that the bounded entity, the χ oscillation, is *constituted* by the curvature drop in the same way the photon is constituted by electroweak symmetry breaking and the axion by PQ breaking plus QCD potential development. The randomness of conventional misalignment is replaced with a Lagrangian determination because the constituting event is now the curvature drop, not the prior selection of a quasi-random initial angle from a flat prior.

Curvature plays a double role in this case: as the geometric curvature R that switches the effective potential from one form to another, and as the field-space curvature $\partial^2 V_{\text{eff}}/\partial\chi^2$ that defines the well in which the post-event oscillations occur. The two senses of curvature converge in this scenario in a way they do not in the photon and axion cases, where field-space curvature is the dominant role. This convergence is part of why the curvature-induced scenario is structurally clarifying: it makes explicit the role spacetime curvature can play as the constituting mechanism, a role obscured in cases where the constituting event is purely internal to the matter sector.

4 The Pattern Across Cases

4.1 What Is Invariant

Across the three case studies, the same three components appear. Each case requires curvature — geometric, potential, or both — to define the well in which the bounded entity sits. Each case requires a symmetry-breaking event in the structural sense — gauge, global, or curvature-mediated — to select a specific configuration from a previously degenerate manifold. Each case requires a scalar (or scalar-like) field to supply the functional content of the bounded entity.

The invariance is not at the level of mechanism. The mechanisms differ wildly. The invariance is at the level of what functional roles must be played for a bounded entity to come into existence at all.

4.2 What Varies

What varies across the cases is which physical structure plays each functional role. In the photon case, curvature is the field-space curvature of $V(\Phi)$; the breaking is a gauge symmetry breaking; the scalar is the Higgs doublet. In the axion case, curvature is the QCD-induced $V(\theta)$ curvature near the minimum; the breaking is two-stage (global $U(1)_{\mathrm{PQ}}$ followed by QCD potential development); the scalar is the axion field. In the curvature-induced scenario, curvature is both spacetime R and the resulting field-space $V_{\mathrm{eff}}(\chi)$; the breaking is the curvature drop at end of inflation; the scalar is χ.

The same template, three different fillings.

4.3 Why These Three

We argue that the three components are jointly necessary, in the following sense.

Without curvature, there is no defined structure within which a bounded entity can be localized. A field with no potential well, or in a region of unstable potential, cannot constitute a stable bounded entity. The photon is bound to its dispersion relation; the axion to its temperature-dependent potential; the χ relic to $V_{\mathrm{eff}}(\chi)$ after the curvature drop. In each case some quantity playing the role of curvature provides the well.

Without symmetry breaking — in the structural sense of an identity-creation event — there is no distinction between the bounded entity and its environment. A field-space well in a fully symmetric configuration admits no entity that is identifiable as distinct from the symmetric phase. The photon as a distinct gauge boson does not exist before electroweak symmetry breaking; the axion as a dark-matter candidate does not exist before QCD potential development; the χ relic does not exist before the curvature drop. The breaking event is what makes "this entity, not that one" a meaningful distinction.

Without a scalar (or scalar-like) field to supply functional content, the bounded entity has nothing to do — no observable effect that would mark it as physically real rather than a mere geometric feature of the theory. The Higgs gives mass; the axion oscillates as dark matter; the χ relic does the same. The functional content is what makes the bounded entity observable as such.

Each component is necessary; the three together are jointly sufficient for the cases we have examined.

4.4 Constituted, Not Triggered

The sharper claim of this paper is that symmetry breaking *constitutes* identity rather than merely triggering it. We have argued for this case-by-case: the photon does not pre-exist electroweak symmetry breaking, the axion-as-dark-matter does not pre-exist QCD potential development, the χ relic does not pre-exist the curvature drop. The claim is general enough to be worth stating in its general form.

The alternative reading would be that the bounded entity exists in some latent or potential form before the breaking, and the breaking merely "activates" it. We do not think this reading survives close inspection. Before electroweak symmetry breaking, there is no field whose properties match those of the photon — no massless gauge boson with the specific charge-coupling pattern Q induces. There are W^a and B fields with different gauge structure, related by a symmetry that mixes them. The photon's identity requires the specific selection $\langle\Phi\rangle$ to fix Q. Without that selection, the photon does not exist in any sense — latent, potential, or actual.

The same is true in the other cases. The axion field exists as a pseudo-Goldstone after PQ breaking, but the axion *as a dark matter candidate* requires the QCD potential development; before that, the field has no preferred minimum to oscillate about and supplies no relic abundance. The χ field exists in the curvature-induced Lagrangian throughout cosmological history, but the χ *relic* — the bounded entity we identify as cold dark matter — exists only after the curvature drop has produced the displacement and the oscillations. In each case, what we name "the entity" is constituted by the breaking event, not merely activated by it.

This claim is partly definitional and partly substantive. Definitional: we define "the entity" as the configuration that exists in the broken phase, so by construction the entity does not exist before the breaking. Substantive: the configuration that exists in the broken phase has features (mass, charge, coupling pattern, oscillation frequency) that are not predicted by the symmetric-phase Lagrangian alone — they are determined by the specific selection the breaking makes. Both senses agree that "constituted" rather than "triggered" is the more accurate description of what the breaking does. The structural realist reading (Ladyman and Ross, 2007; French, 2014) takes this seriously: what is fundamental is the relational pattern by which the entity comes into being, not a prior substantival entity that the breaking modifies.

5 The Constitutive Fingerprint

A structural observation that produces no testable consequence is taxonomy, not physics. We argue that the structural reading offered here produces a falsifiable prediction that generalizes across all three case studies, and that this generalization is the kind of empirical handle by which structural observations earn their physical content. We name this prediction the *constitutive fingerprint* of an identity-creation event.

The prediction is the following. Each identity-creation event leaves three classes of observables, probing the constituting event's *before*, *during*, and *after* phases respectively. The temporal labels are structural shorthand: all observables are measured in the broken phase, but they are classified by which aspect of the constituting event determines their value. Before-phase observables encode features of the symmetric-phase Lagrangian that survive into the broken phase as parameter relations — measurements made now whose values are fixed by pre-event theoretical structure. During-phase observables encode relations that the breaking event itself sets — measurements that probe the symmetry structure of the breaking. After-phase observables encode topological or defect-related features of the post-event configuration. Under the structural reading, these three classes are determined by three different theoretical structures and are therefore functionally independent: there is no parameter of the underlying theory that fixes two of them simultaneously, and no reason within the structural reading for them to correlate beyond what the theory's parameter structure already mandates. A finding that they correlate more strongly than this — that what appear to be three independent fingerprints of a single constituting event are in fact parameters of a single underlying process — would constitute structural evidence against the reading.

The constitutive fingerprint is therefore a three-class observational signature, with each class probing a different temporal phase of the identity-creation event and each determined by a different theoretical structure. We now identify the specific constitutive fingerprints expected in each of the three case studies, and we argue that the same fingerprint pattern recurs across cases that mechanically have nothing in common.

5.1 The Constitutive Fingerprint of Electroweak Symmetry Breaking

The before-phase observable is the Weinberg angle θ_W. Although measured as a parameter of the broken phase, θ_W is determined by the ratio of the gauge couplings of the symmetric-phase Lagrangian: $\tan\theta_W = g'/g$. The angle therefore probes the structure of $SU(2)_L \times U(1)_Y$ prior to the breaking — it is a fingerprint of the symmetric-phase gauge structure that survives into the broken phase as a parameter relation among observable masses and couplings. Its measured value ($\sin^2\theta_W \approx 0.231$ at the Z pole) encodes information about a Lagrangian structure that, on its own, the symmetric phase admits in any of a continuum of configurations.

The during-phase observable is the ρ parameter, $\rho = m_W^2/(m_Z^2 \cos^2\theta_W)$. At tree level, $\rho = 1$ exactly, a relation enforced by the approximate $SU(2)_R$ custodial symmetry of the Higgs sector itself. The ρ parameter probes the symmetry structure of the breaking event — specifically, the fact that the Higgs vev $(1/\sqrt{2})(0, v)$ is invariant under custodial $SU(2)_R$ as well as under the gauge $U(1)_{\text{em}}$. Loop corrections shift ρ by parts in 10^4, and precision electroweak measurements of ρ probe the breaking event with corresponding sensitivity. This is a fingerprint of the breaking itself, not of the prior gauge structure or of post-event topology.

The after-phase observable is the topology of the broken vacuum manifold. For the Standard Model breaking $SU(2)_L \times U(1)_Y \to U(1)_{\text{em}}$, the relevant homotopy groups π_n of the vacuum manifold are trivial for n that would generate cosmologically relevant defects: no electroweak-scale magnetic monopoles arise from this breaking (those would require a higher-rank breaking such as a grand unification), no cosmic strings, no domain walls (Kolb and Turner, 1990). The observed absence of electroweak-scale topological defects in our universe is a positive after-phase observable — the topology of the post-event configuration is empirically constrained by what we do not see.

The constitutive fingerprint of electroweak symmetry breaking is therefore the triple $\{\theta_W, \rho, \text{topological triviality}\}$. Under the structural reading, these three observables are determined by three different theoretical structures: θ_W by the symmetric-phase coupling ratio, ρ by the custodial symmetry of the breaking event, topological triviality by the homotopy structure of the broken vacuum manifold. The reading predicts that they are functionally independent. Detection of an unexpected correlation — for example, a deviation from $\rho = 1$ that tracked changes in θ_W in a way not predicted by Standard Model loop corrections, or a topological defect signature whose presence covaried with precision-electroweak parameters — would suggest that what we identified as three independent aspects of EWSB is in fact a parameter of a deeper underlying process. The current empirical situation is consistent with independence, and the structural reading predicts that this independence is not an accident.

5.2 The Constitutive Fingerprint of the QCD Axion's Stacked Breakings

The before-phase observable is the isocurvature spectrum imprinted on cold dark matter by inflationary fluctuations of the axion field. In the pre-inflationary PQ-breaking case, the axion field undergoes quantum fluctuations of order $\delta\theta \sim H_I/(2\pi f_a)$ during inflation, which propagate into the post-event isocurvature power spectrum (Kolb and Turner, 1990; Planck Collaboration, 2020). The amplitude is sensitive to the inflationary scale H_I and to the decay constant f_a. Both quantities characterize the configuration prior to the QCD-scale event that constitutes the axion as a dark matter candidate; the isocurvature signature is therefore a before-phase fingerprint of the axion's identity-creation event.

The during-phase observable is the relationship between the axion mass and decay constant set by the QCD-induced potential: $m_a \times f_a \approx \Lambda_{\text{QCD}}^2$ (more precisely, $m_a^2 f_a^2 \approx m_\pi^2 f_\pi^2 \cdot m_u m_d/(m_u + m_d)^2$). This relation is determined by the breaking event itself — specifically, by the QCD instanton effects that develop the axion potential at temperatures near Λ_{QCD}. The mass scale of the resulting dark matter candidate and its decay constant are not independent parameters, but locked together by the structure of the constituting event. Any axion detection that simultaneously measures m_a and f_a probes the during-phase event directly.

The after-phase observable is the cosmological status of axion topological defects. In post-inflationary PQ-breaking scenarios, the field takes random values across causally disconnected patches, producing axion strings and (after QCD potential development) domain walls bounded by those strings. The wall network is a cosmological catastrophe unless the domain wall number $N_{\text{DW}} = 1$, where N_{DW} is a topological invariant set by the structure of the QCD-induced potential. The observation that our universe does *not* contain a cosmologically dominant domain wall network constrains the topology of the post-event configuration directly — and through it, constrains either the order of the stacked breakings (pre- vs post-inflationary PQ) or the value of N_{DW} in nature.

The constitutive fingerprint of the QCD axion's stacked breakings is the triple {isocurvature amplitude, m_a–f_a relation, domain wall topology}. These three observables are determined by three different theoretical structures: the isocurvature amplitude by H_I and f_a; the $m_a \times f_a$ relation by Λ_{QCD}; the domain wall problem by N_{DW}. They are functionally independent under the structural reading. A finding that they correlated in unexpected ways — for example, that observed isocurvature constraints implied $m_a \times f_a$ values incompatible with QCD-induced expectations — would challenge the reading that PQ breaking and QCD potential development are two distinct identity-creation events rather than parameters of a single underlying process.

5.3 The Constitutive Fingerprint of the Curvature-Induced Scenario

The before-phase observable is the suppressed isocurvature spectrum predicted by the curvature-induced scenario. During inflation, χ is held at the minimum of an effective potential with curvature $m_{\text{eff}}^2 \approx 12\xi H_I^2$. Quantum fluctuations about this minimum scale as H_I/m_{eff}, suppressed relative to the canonical misalignment case by a factor $1/\sqrt{12\xi}$ (Markkanen et al., 2018). The amplitude of these fluctuations is determined by H_I and ξ — both pre-event quantities (ξ is a Lagrangian parameter; H_I characterizes the inflationary epoch).

The during-phase observable is the H_I-independence of the relic abundance. In the regime $\xi H_I^2 \gg m^2$, the inflationary minimum sits at $\chi_{\text{inf}} \approx m^2 v/(12\xi H_I^2) \ll v$, and the post-event displacement is $\Delta\chi \approx v$, independent of H_I. The relic abundance therefore depends on m and v but not on H_I above threshold. The H_I-independence is a fingerprint of the curvature-drop event itself: the constituting event is the transition from $R \approx 12H_I^2$ to $R \approx 0$, and the post-event displacement is determined by the geometry of this transition rather than by the inflationary scale. A future tensor-to-scalar measurement that fixes H_I (by, for example, LiteBIRD or CMB-S4) tests this prediction directly: it should find no correlation between H_I and the χ relic abundance once H_I exceeds the threshold.

The after-phase observable is the absence of domain walls. Because v is a fixed Lagrangian parameter rather than a randomly broken $\mathbb{Z}_2$ direction, the post-event minimum is unique. No domain wall network forms. This is a topological feature of the post-event configuration, set by the uniqueness of v in the Lagrangian.

The constitutive fingerprint of the curvature-induced scenario is the triple {suppressed isocurvature, H_I-independence, no domain walls}. The three observables are determined by three different theoretical structures: the isocurvature suppression by H_I and ξ; the H_I-independence by the regime $\xi H_I^2 \gg m^2$ and the geometry of the curvature drop; the absence of domain walls by the uniqueness of v. They are functionally independent under the structural reading. A finding that they correlated — for example, that the predicted isocurvature suppression depended on whether domain walls were observed in particular patches — would challenge the reading that the curvature drop is a single identity-creation event.

5.4 What Generalizes

The structural reading makes the same prediction across all three cases. Each identity-creation event produces a constitutive fingerprint: three classes of observables, determined by three different theoretical structures, predicted to be functionally independent. The specific observables differ — θ_W versus isocurvature amplitude versus suppressed isocurvature; ρ versus the m_a–f_a relation versus H_I-insensitivity; topological triviality versus the domain wall problem versus the uniqueness of v — but the constitutive-fingerprint pattern of three independent fingerprints recurs.

This generalization is what gives the structural reading empirical traction beyond any single case. The reading is not merely a way of organizing the existing predictions of one mechanism; it makes the same kind of independence prediction across every case where it claims to apply. The constitutive fingerprint is therefore testable in three distinct empirical regimes, with three different sets of observables, by three different communities of physicists. Each test of any one of them probes the structural reading.

It might be objected that the constitutive-fingerprint prediction is weaker than what individual mechanical theories provide. Each case's underlying mechanism predicts specific quantitative values for each observable; the structural reading predicts only that the three are independent. This comparison, however, mistakes the level of analysis. The structural reading is not a competing mechanical theory operating on the same quantitative-prediction metric. It is a hypothesis about pattern recurrence across cases that no mechanical theory — by its own design — can make. No mechanical theory of electroweak symmetry breaking predicts that the QCD axion or the curvature-induced relic should exhibit the same independence structure; mechanical theories operate within their respective sectors. The structural reading predicts that the same constitutive-fingerprint pattern recurs across mechanically distinct sectors, and this prediction has empirical traction of its own.

Pattern recurrence across mechanically distinct cases is the kind of evidence that distinguishes structural-realist hypotheses from purely mechanical alternatives (Ladyman and Ross, 2007; French, 2014). A structural realist holds that what is fundamental in physics is not the inventory of entities but the relational and structural patterns that recur across them. Empirically supporting such a position requires showing that the same patterns appear in multiple cases that have no shared mechanism. That is precisely what the constitutive-fingerprint prediction provides: a structural pattern that recurs across cases with nothing else in common. A reading that organizes a single case is a reframing; a reading that makes the same predictive structure visible across mechanically distinct cases is a hypothesis about how physical reality is organized — and pattern recurrence is the form of empirical support such hypotheses can earn.

We do not here calculate the expected magnitude of the correlations under mechanical-variant null hypotheses, nor do we identify the specific cosmological observations and precision-electroweak measurements that would test the constitutive-fingerprint predictions with sufficient sensitivity. These are quantitative tasks the present paper does not undertake. What we do claim is that the structural reading produces a generalizable falsifiable prediction — and that the prediction's structural recurrence across mechanically distinct cases is itself empirical evidence (though not yet decisive evidence) that the structural reading captures something real about how bounded physical entities come into being.

6 Honest Limitations

The structural reading proposed here has limitations that we state directly.

The framework is observational, not predictive of new physics. The structural template does not predict the existence of any new particle, field, or interaction. It claims to recognize a pattern in the way bounded physical entities come into being, and to use that pattern to clarify the relations among existing scenarios. A reader who treats the framework as a substitute for mechanical model-building will be disappointed. The framework's role is to make explicit the structural commitments that mechanical models make tacitly, and to predict that those commitments leave a recurring empirical signature in the form of a constitutive fingerprint.

The "scalar field" component is loose. We have written as if scalar fields are the only carriers of the functional content of bounded entities, but this is not strictly true. Fermion bilinears can play scalar-like roles (the QCD chiral condensate $\langle \bar{q}q \rangle$ is a paradigmatic case). In some contexts, vector boson condensates or composite operators play the same functional role. We have used "scalar (or scalar-like)" throughout to flag this, but the looseness is a real feature of the framework. A more rigorous treatment would specify the conditions under which a non-scalar object can carry the functional content of a bounded entity, and we have not provided that specification.

Not every physical entity obviously fits the template. The graviton, if it exists as a fundamental quantum, is not produced by any symmetry-breaking event we have identified. Standard Model fermions are constituted partly by electroweak symmetry breaking (which gives them mass) but also by their gauge structure, which exists prior to any breaking. We do not claim the structural template is universal across all of physics; we claim it captures the constituting pattern of the bounded entities examined here, and we conjecture that it extends to a broader class of cases without claiming to have established that extension.

The "constituted, not triggered" distinction depends on identity assignment. If one considers the photon to be "already there" in some abstract sense before electroweak symmetry breaking — present in the gauge structure of $SU(2)_L \times U(1)_Y$ as a possible linear combination — then the breaking merely makes that pre-existing combination the privileged one. We have argued against this reading on the grounds that the labels "photon," "Z," "θ_W" are themselves identified by the breaking, and a structural realist reading (Ladyman and Ross, 2007) supports our position by treating the relational structure as ontologically prior. But the argument depends on a specific assignment of identity. Readers committed to a different assignment may disagree with the constitutive claim while accepting the rest of the framework, including the constitutive-fingerprint prediction.

The constitutive-fingerprint prediction is structural rather than quantitative. The structural reading does not predict specific numerical values for the case-by-case observables; it predicts a pattern of functional independence within each case and the recurrence of that pattern across mechanically distinct cases. This is empirical content, but of a different kind than a mechanical theory's case-specific numerical predictions. A reader expecting the structural reading to compete with mechanical theories on numerical-prediction grounds will find them addressing different questions: the mechanical theory provides numerical specificity within a single case, while the structural reading provides pattern recurrence across cases. The contribution is at a different level of analysis, not a deficiency on the same one.

Engagement with philosophical work on identity in physics is limited. There is a substantial literature on identity, individuation, and emergence in philosophy of physics — including the structural realism literature (Ladyman and Ross, 2007; French, 2014; Esfeld, 2004) — where related questions are discussed at greater length and with more conceptual care than the present paper offers. The structural framing here is meant to be operational rather than philosophically rigorous. A reader interested in the philosophical implications of the constitutive claim should consult that literature directly rather than treating the present paper as a substitute for it.

These limitations do not, in our view, undermine the structural reading. They identify where the work remains to be done: a more rigorous treatment of scalar-content's loose membership; an explicit examination of cases (gravitons, fermions) that may or may not fit the template; deeper engagement with the structural realism literature; quantitative development of the constitutive-fingerprint prediction with calibrated null hypotheses for each case. Each is a direction in which the framework could be sharpened.

7 Future Directions

The structural pattern identified here — bounded physical entities constituted by symmetry breaking within a curvature-defined potential, with scalar field dynamics supplying functional content — has analogs outside high-energy physics that we have deliberately not pursued in the body of this paper. We sketch them briefly, on the principle that gesturing at the broader resonance is appropriate while developing it would require substantially more work than the present synthesis warrants.

In condensed matter physics, phase transitions accompanied by spontaneous symmetry breaking constitute new bounded entities — Cooper pairs in superconductors, magnetic domains in ferromagnets, ordered states in liquid crystals — whose identity arises from the breaking event. The Landau–Ginzburg framework formalizes the order parameter (a scalar or scalar-like quantity) whose dynamics supply the broken phase's functional content. The triadic template appears applicable, with curvature played by the Landau free-energy potential, breaking played by the phase transition, and scalar content played by the order parameter. The constitutive-fingerprint prediction should also have analogs in condensed-matter systems: pre-transition susceptibilities, transition-temperature relations, and post-transition topological defects (vortices in superfluids, domain walls in ferromagnets) are determined by different theoretical structures and should be functionally independent. Whether the constitutive claim transfers to this domain — whether, for example, the superconducting state genuinely constitutes a new bounded entity rather than activating a latent one — is an open question we leave to that literature.

In developmental biology, cell-fate determination through gene-regulatory dynamics has been formalized using Waddington's epigenetic landscape, with attractor states that cells settle into through a symmetry-breaking selection of one developmental trajectory from an initially degenerate set. Concentration thresholds and morphogen gradients play roles structurally analogous to the curvature and breaking components. We do not claim the constitutive-fingerprint pattern applies in any rigorous sense; we note only that the broader triadic structure recurs at a level of generality that suggests it is not unique to particle physics.

Other potential domains include the formation of stable self-organizing structures in non-equilibrium thermodynamics, where dissipative structures emerge through symmetry-breaking transitions in driven systems; the emergence of distinct linguistic identities through the breaking of phonetic continua under social and articulatory constraints; and the constitution of coherent organizational entities in multi-agent systems through the breaking of behavioral degeneracy. The cross-domain extension of triadic structural patterns has been explored under the broader framework of Morales (2026), to which we refer the interested reader. The present paper restricts attention to the high-energy physics instance and the empirical content the constitutive-fingerprint prediction produces.

The most rigorous next step within high-energy physics would be a quantitative development of the constitutive-fingerprint prediction. A calculation of the expected correlation structure under mechanical-variant null hypotheses, combined with a sensitivity analysis for current and proposed observations — Planck Collaboration (2020), LiteBIRD, CMB-S4 for the dark matter cases; precision electroweak data from LEP and the LHC for EWSB — would convert the structural reading from an interpretive framework into a sharply testable hypothesis. We commend this calculation to the technical literature.

8 Conclusion

We have argued that bounded physical entities — stable particles, cosmological relics, broken-symmetry vacuum states — are constituted by the joint operation of three structural components: a curvature that defines a potential well, a symmetry-breaking event that selects a specific configuration from a previously degenerate manifold, and a scalar field that supplies the entity's functional content. The mechanical machinery underlying these components differs widely across cases. The structural pattern by which the three components combine to constitute a bounded entity does not.

The sharper claim, developed across the case studies, is that symmetry breaking is *constitutive* rather than merely triggering. The photon does not pre-exist electroweak symmetry breaking. The QCD axion as a dark matter candidate does not pre-exist the development of its temperature-dependent potential. The χ relic does not pre-exist the curvature drop at the end of inflation. In each case, what we name the entity — the configuration with specific mass, charge, coupling pattern, oscillation frequency — is what the breaking event produces. The labels by which we identify the bounded entity are themselves selected by the breaking. Identity, in this strong sense, is what symmetry breaking constitutes.

The structural reading earns its physical content through the constitutive-fingerprint prediction. Each identity-creation event produces three classes of observables — features encoded in the parameter relations of the symmetric phase, relations set by the breaking event, and topological features of the post-event configuration — that are determined by three different theoretical structures and predicted to be functionally independent. The specific fingerprints differ across cases. For electroweak symmetry breaking, they are the Weinberg angle, the ρ parameter, and topological triviality of the broken vacuum manifold. For the QCD axion, they are the isocurvature amplitude, the m_a–f_a relation, and the domain wall problem. For the curvature-induced misalignment scenario, they are the suppressed isocurvature spectrum, the H_I-independence of the relic abundance, and the absence of domain walls. The constitutive-fingerprint pattern recurs across mechanically distinct cases, and the recurrence is the structural reading's principal empirical content.

We have stated the framework's limitations openly. It is observational rather than predictive of new physics. The scalar-content component is loose. Not every physical entity obviously fits the template — gravitons, fermions, and other cases may require modification or extension. The constitutive-versus-triggering distinction depends on how one assigns identity. The constitutive-fingerprint prediction is structural rather than quantitative. Engagement with the philosophical literature on identity in physics is limited.

What the structural reading offers, despite these limitations, is a vocabulary for asking the structural question precisely. The question of how anything physical comes to be a bounded *something* — distinct, identifiable, with its own functional content — has, across the cases we examined, a structurally common answer. Curvature sets the well. Symmetry breaking constitutes the identity. Scalar content supplies the function. Each component is necessary; the three together are jointly sufficient for the cases at hand. And the empirical fingerprint of this structural pattern — three independent observables probing the constituting event's before, during, and after phases — provides a falsifiable test that recurs in the same form across mechanically distinct sectors of physics. That recurrence is what makes the pattern more than taxonomy. It is what makes it a hypothesis about how bounded physical reality is organized.

References

Abbott, L. F., and P. Sikivie. 1983. “A Cosmological Bound on the Invisible Axion.” *Physics Letters B* 120: 133–136.

ATLAS Collaboration. 2012. “Observation of a New Particle in the Search for the Standard Model Higgs Boson with the ATLAS Detector at the LHC.” *Physics Letters B* 716: 1–29.

Bezrukov, F., and M. Shaposhnikov. 2008. “The Standard Model Higgs Boson as the Inflaton.” *Physics Letters B* 659: 703–706.

CMS Collaboration. 2012. “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC.” *Physics Letters B* 716: 30–61.

Dine, M., and W. Fischler. 1983. “The Not-So-Harmless Axion.” *Physics Letters B* 120: 137–141.

Esfeld, M. 2004. “Quantum Entanglement and a Metaphysics of Relations.” *Studies in History and Philosophy of Modern Physics* 35: 601–617.

French, S. 2014. *The Structure of the World: Metaphysics and Representation.* Oxford: Oxford University Press.

Glashow, S. L. 1961. "Partial-symmetries of Weak Interactions." *Nuclear Physics* 22: 579–588.

Kim, J. E. 1979. "Weak-Interaction Singlet and Strong CP Invariance." *Physical Review Letters* 43: 103–107.

Kolb, E. W., and M. S. Turner. 1990. *The Early Universe.* Reading, MA: Addison-Wesley.

Ladyman, J., and D. Ross. 2007. *Every Thing Must Go: Metaphysics Naturalized.* Oxford: Oxford University Press.

Markkanen, T., A. Rajantie, and S. Stopyra. 2018. "Cosmological Aspects of Higgs Vacuum Metastability." *Frontiers in Astronomy and Space Sciences* 5: 40.

Morales, J. R. 2026. *Triadic Constraint and Emergent Stability: A Meta-Theoretical Framework in Philosophy of Physics — The Singularity–Duality–Triality Framework.* Manuscript. PhilPapers.

Peccei, R. D., and H. R. Quinn. 1977. "CP Conservation in the Presence of Pseudoparticles." *Physical Review Letters* 38: 1440–1443.

Planck Collaboration. 2020. "Planck 2018 Results. VI. Cosmological Parameters." *Astronomy & Astrophysics* 641: A6.

Preskill, J., M. B. Wise, and F. Wilczek. 1983. "Cosmology of the Invisible Axion." *Physics Letters B* 120: 127–132.

Salam, A. 1968. "Weak and Electromagnetic Interactions." In *Elementary Particle Theory: Relativistic Groups and Analyticity*, edited by N. Svartholm, 367–377. Stockholm: Almqvist & Wiksell.

Shifman, M. A., A. I. Vainshtein, and V. I. Zakharov. 1980. "Can Confinement Ensure Natural CP Invariance of Strong Interactions?" *Nuclear Physics B* 166: 493–506.

Tenkanen, T. 2019. "Dark Matter from Scalar Field Fluctuations." *Physical Review Letters* 123: 061302.

Weinberg, S. 1967. "A Model of Leptons." *Physical Review Letters* 19: 1264–1266.