

The Geometry of the Moral Law

Santiago De Jesus Sanchez Borboa

Abstract

This paper builds on the framework introduced in “Symmetry and the Form of Ethical Objectivity.” There, I argue that objectivity in practical evaluation is appropriately understood as invariance under transformations that preserve all normatively relevant structure, and that once even minimal invariance constraints are accepted, symmetry requirements bootstrap from diachronic coherence to robust impartiality. The present paper extends that argument from global to local invariance and draws out its implications for the foundations of moral philosophy. When the demand for objectivity is strengthened from global symmetries (uniform transformations of time, person, and perspective) to local invariance (frame-independence under evaluative recalibrations that may vary from agent to agent and context to context), the mathematics of gauge theory shows that a new dynamical structure is necessitated: a connection field that compensates for local variation and enables coherent comparison of evaluative judgments across differing deliberative frames. I argue that this formally derived structure has precisely the properties that Kant attributes to the good will. Its value is unconditional because it is a gauge-invariant feature of the evaluative structure itself, not a contingent configuration of ethical content. It is the condition of possibility for frame-independent ethical evaluation, not one more object of evaluation among others. The moral law, on this account, is a local invariance requirement on practical reason; the good will is the connection field that requirement generates. The paper develops the formal apparatus, works through the central philosophical implications, and examines where the analogy between gauge theory and moral philosophy breaks down, showing them to be philosophically productive rather than problematic.

Introduction

Kant claims that the good will is the only thing good “without limitation”, that it “shines by its own light, like a jewel,” independently of what it achieves or fails to achieve.¹ The claim is among the most celebrated in moral philosophy. It is also among the least well understood. What does it mean for a practical orientation to have unconditional value? What structural feature of the good will could make it the case that its worth does not depend on the agent’s circumstances, the consequences of action, or the evaluative standpoint from which it is assessed? Kant’s own answer, namely that the good will is determined by the moral law rather than by inclination, relocates the question without fully resolving it. For the moral law itself is articulated through the categorical imperative, and the categorical imperative is applied by agents embedded in particular evaluative situations, with particular vocabularies, scales, and ways of registering what matters. If the good will’s value is truly unconditional, there should be a precise sense in which it is independent of all such local variation. The present paper offers formal tools for making this independence precise. Drawing on the structural logic of *gauge theory* (a framework that, in physics, characterizes what remains objective when local representational conventions are allowed to vary freely), I argue that the unconditional value of the good will can be given a precise formal characterization as a field with *gauge-invariant content*, i.e., as evaluative structure that survives every local transformation of agents’ representational frames. The good will “shines by its own light” because its contribution to the ethical landscape is constitutively independent of the local calibrations through which any particular agent describes it.

This claim may invite suspicion. Ethics is not physics; moral agents are not charged particles; the categorical imperative is not a field equation. The paper takes these differences seriously (§4 is devoted to examining the disanalogies and arguing that they are productive rather than problematic). The claim is not that morality *is* a gauge theory in any literal sense. It is that the structural logic of local invariance—the logic that asks what remains objective when representational conventions are allowed to vary from point to point—applies to the ethical domain with genuine formal precision, and that when it is applied, it yields results that illuminate the Kantian framework in ways not previously available.

Plan of the paper

Section 1 recapitulates the companion paper’s framework and motivates the transition from global to local invariance: the key philosophical move on which the rest of the paper depends.

Section 2 develops the formal apparatus. Working within the discrete framework of lattice gauge theory (the natural setting for a domain of finitely many agents), it introduces local evaluative frames and local gauge transformations,

¹Kant, *Groundwork of the Metaphysics of Morals*, 4:393–394. All references to Kant follow the standard Akademie edition pagination.

derives the ethical connection field as the structure necessitated by ethical local invariance, establishes its transformation law, and characterizes this field's curvature as the gauge-invariant content. Two running examples, one involving numerical scaling of evaluative magnitudes (an abelian gauge group) and one involving cross-cultural conceptual vocabularies (a nonabelian gauge group), are developed throughout.

Section 3 argues that the good will fits the structural profile of the ethical connection field (i.e., it enforces coherence across locally varying frames, is necessitated by the demand for practical objectivity, varies in its local expression, and generates gauge-invariant content through its curvature) and that its unconditional value corresponds to its gauge-invariant character.

Section 4 examines the principal disanalogies between the physical and ethical cases: the non-uniqueness of the ethical connection field, the self-referential character of this field, the absence of a variational principle, and the normative (rather than descriptive) character of the ethical connection. It argues that these disanalogies are philosophically productive and don't undermine the larger structural analogy between the physical and ethical domain spelled out by the paper. Rather, the disanalogies constitute places where the target domain (ethics) has distinctive structure that the source domain (gauge theory) doesn't capture but does help us see.

Section 5 concludes by noting the methodological parallel between the Kant and Einstein and by flagging the framework's resources for dissolving the longstanding "empty formalism" objection to Kantian ethics as directions for future investigation.

The argument is conditional, as befits a structural investigation. It does not presuppose that Kant's ethics is correct; it shows what formal structure the ethical domain must have *if* OI is accepted and local invariance is imposed, and it argues that the resulting structure illuminates the Kantian framework with a precision not previously available. Readers who do not share Kantian commitments may nonetheless find the formal results of independent interest: the derivation of the ethical connection field from an ethical local invariance requirement, the characterization of this field's curvature as gauge-invariant content, and the discussion of disanalogies between the physical and ethical cases are all developed at a level of generality that does not depend on the identification with the good will.

1. From Global to Local Invariance

1.1 The Invariance Framework Recapitulated

A companion paper develops objectivity-as-invariance (OI) as a bridge principle connecting the structure of practical reason to substantive ethical constraints.²

²Sanchez Borboa, "Symmetry and the Form of Ethical Objectivity," (manuscript A). The present section summarizes only what is needed to motivate the extension; the reader is referred to the companion paper for the full argument and formal apparatus.

The guiding idea is drawn from the methodology of modern physics: a feature of an evaluation counts as objective just in case it is invariant under transformations that alter what is normatively irrelevant. Features that shift under such transformations are artifacts of the evaluator’s representational standpoint, not of the subject matter.

OI is not itself a moral claim. It does not tell us which transformations are normatively irrelevant; it tells us what follows once a class of transformations is so identified. In this respect it functions like a structural constraint, a condition on the form an objective evaluation must take, not on its content.

The companion paper argues that once OI is accepted as a constraint on practical reason, a bootstrapping dynamic emerges. The most minimal invariance requirement available to a deliberating agent is diachronic coherence: the evaluation should not depend on the bare temporal index at which it is issued. This is invariance under time-translation within a single deliberative life, and it is difficult to reject without undermining the very activity of sustained deliberation.

But diachronic coherence is already an instance of OI: it treats the bare temporal index as normatively irrelevant. And once the agent is in the business of identifying features of the evaluative standpoint as irrelevant, a principled question arises about every other candidate feature. The companion paper argues (via a parity argument and a diagnostic analysis of the logical structure of egoism) that once one accepts invariance under temporal translation, it is arbitrary not to accept invariance under permutation of agents. The bare indexical fact that an evaluation is issued by me rather than by you is, from the standpoint of OI, structurally analogous to the bare indexical fact that it is issued now rather than then. If the latter is normatively irrelevant, so is the former.

The result is robust ethical impartiality: objective evaluations must be invariant under permutations of the agents who occupy evaluative roles. Formally, if a symmetry group G acts on the space of agent-indexed evaluative standpoints, then an objective evaluation E must factor through the quotient map q that identifies standpoints related by the group action ($E = \overline{E} \circ q$) so that only the G -invariant content of the standpoint is registered by the evaluation.³

This is a powerful result. It delivers impartiality not as a first-order moral premise but as a structural consequence of taking the aspiration to objectivity seriously and recognizing that bare indexical features of the evaluative standpoint are representational artifacts. But it is also, in an important respect, incomplete.

³The factorization condition is developed precisely in the formal appendix of the companion paper, where it is used to characterize the sense in which prudential reasons are conditional on the evaluator’s identity: they are invariant under time-translation but not under agent-permutation, and this asymmetry is made visible by the failure of factorization through the larger quotient.

1.2 The Limitation of Global Invariance

The invariances the companion paper focuses on are global: the same transformation (e.g., the same permutation of agent labels) is applied uniformly across all agents simultaneously. This is the ethical analog of what physicists call a global symmetry: a transformation that acts identically at every point of the relevant domain.

Global symmetries are far from trivial. In physics, Noether’s first theorem shows that each continuous global symmetry of a system’s dynamics corresponds to a conserved quantity: translational symmetry yields conservation of momentum, rotational symmetry yields conservation of angular momentum, and so on (Noether, 1918/2011). Analogously, global agent-permutation invariance in the ethical case yields a conserved evaluative structure, impartiality, that constrains the form objective evaluations can take.

But global symmetry, for all its power, rests on a tacit idealization. To demand that evaluations be invariant under a uniform relabeling of agents is to presuppose that all agents are, as it were, already described within a single shared evaluative frame. The permutation swaps who occupies which role, but it assumes that the roles themselves (i.e., the way the evaluative landscape is parsed and the vocabulary in which ethically relevant features are registered) are held fixed across agents.

This idealization is substantial, and once noticed, it is difficult to regard as innocent. Consider: two agents may assign different numerical weights to the same ethically relevant feature, one calibrating suffering on a scale where moderate pain registers as 3, the other on a scale where it registers as 7. If these are genuinely different scales rather than different judgments (if, that is, the disagreement is about the unit of measurement rather than about the ethical significance of suffering), then the difference is a representational artifact, exactly the kind of thing OI says objective evaluation should be insensitive to. But global invariance does not address it. A uniform permutation swaps agent labels; it does not touch the local calibration each agent brings to the evaluative task.

Or consider agents embedded in different cultural-linguistic traditions who carve up the ethically relevant features of a situation using different conceptual vocabularies, different ways of categorizing obligations, relationships, or ethically salient characteristics of actions. If two traditions recognize the same underlying ethical structure but articulate it through different local terminologies, then the differences between their articulations are, from OI’s standpoint, representational. But a global permutation, which merely asks whether swapping the agents’ positions changes the evaluation, has nothing to say about whether the evaluation is robust across the agents’ different ways of describing their evaluative situations.

The pattern here is general. Global invariance asks: does it matter who occupies a given evaluative role? Local invariance asks a deeper question: does it matter how each agent locally represents the evaluative role they occupy? The first filters out the bare indexical “I” as a representational artifact; the

second filters out the residual frame-dependence that remains even after the “I” has been removed.

If OI is the right bridge principle, i.e., if objectivity really is invariance under transformations that alter only what is normatively irrelevant, then there is principled pressure to move from global to local invariance. The same diagnostic that revealed bare agent-identity as a representational artifact also reveals local evaluative calibration as a representational artifact, at least insofar as the calibration concerns solely the scale or vocabulary of evaluation rather than its substantive content. And once a feature is identified as artifactual, allowing the evaluation to depend on it is, by OI’s own lights, a defect.

1.3 What Local Invariance Demands

The transition from global to local invariance is not a minor technical refinement. In physics, it is the transition from conservation laws to gauge theory, from constraints on what is preserved to constraints on the structure that governs change. Global symmetries stabilize quantities; local symmetries, because they demand invariance under transformations that vary from point to point, force the introduction of new structure — a connection field — that tracks how locally varying descriptions are to be compared (Noether 1918/2011). The connection is not optional; it is the minimum additional apparatus needed for the theory to remain coherent once local invariance is imposed. And the curvature of the connection, i.e., the extent to which the connection’s own structure is nontrivial, constitutes the gauge-invariant dynamical content that emerges from the local symmetry requirement. It is, in a precise sense, the content that does not depend on any local choice of frame.

The guiding conjecture of the present paper is that the ethical case exhibits the same structural pattern. If objective ethical evaluation must be invariant not merely under global transformations (like agent-permutation) but under locally varying re-framings of agents’ evaluative standpoints, then the framework requires additional structure — an ethical analog of the connection field — that enforces coherence across locally varying frames. The central task of §2 is to develop this structure formally. The central claim of §3 is that the structure so derived can be identified with what Kant calls the good will, and that its unconditional value, its “shining by its own light” independent of what it accomplishes, is a consequence of its gauge-invariant character: it is precisely the evaluative content that survives every local re-framing. Before turning to the formal development, it is worth pausing to note what this approach does and does not promise. It does not promise to derive the full content of substantive first-order ethics from the bare requirement of ethical local invariance, any more than gauge theory derives the specific forces of nature from the bare requirement of local symmetry. But crucially, it does not merely constrain the form ethics must take while leaving that form contentless. Just as the gauge principle, once imposed, necessitates the existence of a connection field with determinate structural properties and nontrivial dynamical content (curvature), so too the requirement of local evaluative invariance necessitates the existence

of a structure with determinate properties that does real normative work, as §2 will argue. The formal requirement generates genuine content; what it does not do is fix that content uniquely or exhaustively. Establishing what the requirement demands (and showing that something must play the structural role of the connection) is itself a substantive philosophical achievement, not a merely formal one. The disanalogies between the physical and ethical cases, including the question of what determines the remaining content, are discussed in §4, where they prove to be at least as illuminating as the analogies.

2. The Formal Apparatus: Ethical Local Invariance, Connection Field, and Curvature

The previous section motivated the transition from global to local invariance: if objectivity-as-invariance (OI) requires that evaluation be insensitive to representational artifacts, and if agents' local evaluative frames (i.e., their scales, vocabularies, and calibrations) can vary independently, then a residual frame-dependence persists even after global symmetries have been imposed. The present section develops the formal apparatus that local invariance requires. The argument has three stages:

1. First, I set up the formal landscape and state the ethical local invariance requirement precisely (§2.1–2.2).
2. Then, I show that this requirement cannot be satisfied without additional structure (an *ethical connection field*) and derive the transformation law it must obey (§2.3).
3. Finally, I characterize the *curvature* or *field strength* of the connection field as the genuinely gauge-invariant content: the evaluative content that survives every local re-framing (§2.4).

A note on method is in order before beginning the argument. The formalism presented here works with a discrete set of agents rather than the continuous spacetime manifold of physics. This is not a simplification introduced for the sake of accessibility (though it has that benefit); it is the natural setting for the ethical application, where the relevant domain consists of distinct ethical agents rather than a continuum of spacetime points. The discrete framework has independent mathematical credentials: it is the setting of *lattice gauge theory*, which is both foundational and rigorous.⁴ The structural logic—local invariance forces a connection; curvature characterizes gauge-invariant content—is identical in the discrete and continuous cases. What changes are the analytical tools, not the conceptual architecture.

⁴For the mathematical framework of discrete (lattice) gauge theory, see, e.g., Wilson (1974) and Kogut (1979). The key point for present purposes is that the derivation of the connection from the local invariance requirement does not depend on differential-geometric machinery; it depends on the *logic* of compensating for locally varying transformations, which is the same whether the underlying domain is discrete or continuous.

2.1 The Formal Landscape

Let's begin by recalling the formal vocabulary introduced in the companion paper's appendix, extending it only where the move to local invariance requires.⁵

2.1.1 Cases, evaluations, and evaluative outputs

As in the companion paper, we can fix a domain W of ethically relevant **cases**: the actions, attitudes, maxims, policies, social arrangements, etc. that are the objects of ethical evaluation. A case $w \in W$ is a complete specification of an ethically evaluable situation.

An **evaluation** is a function

$$E : W \longrightarrow V$$

that assigns to each case $w \in W$ an **evaluative output** $E(w) \in V$. The set V of possible evaluative outputs is deliberately left general: it might consist of simple verdicts (permissible, impermissible), richer rankings, or weighted reason-structures. What matters is that E assigns an evaluative status to each case, and that we can ask whether this assignment is sensitive to features it should not be sensitive to.

2.1.2 Agents and local evaluative frames

Let I be a finite set of **agents**, i.e., the ethical subjects whose evaluative standpoints are relevant to the assessment. Each agent $i \in I$ occupies an evaluative standpoint: they register the ethical relevant features of a case and form evaluative judgments on the basis of what they register. The companion paper's global invariance result established that objective evaluations must be invariant under permutations of I , i.e., under relabeling which agent occupies which standpoint.

We can now introduce the additional structure that §1 motivated. Each agent $i \in I$ has a **local evaluative frame** f_i : a calibration that determines how ethically relevant features are represented from i 's standpoint. The frame encompasses:

- the *scales* on which ethically relevant magnitudes are measured (how numerically “large” a given degree of suffering registers);
- the *conceptual vocabulary* through which ethically relevant distinctions are articulated (which categories are used to organize obligations, relationships, or ethically salient features of actions); and
- more abstractly, any *representational parameter* that can vary across agents without, in itself, constituting a substantive ethical disagreement.

⁵Readers who have worked through the companion paper's appendix will find the notation intentionally consistent. The extension consists in the introduction of local evaluative frames (§2.1.2) and local gauge transformations (§2.2), which have no counterpart in the global setting.

The crucial distinction, introduced informally in §1 and now made precise, is between *variation in frame* and *variation in judgment*. Two agents who assign different weights to suffering because they *disagree about how much suffering matters* differ in ethical judgment. Two agents who assign different numerical values to the same degree of suffering because they *calibrate their scales differently* differ in frame. OI is concerned with the latter: frame-variation is representational, and objective evaluation should be insensitive to it.

2.1.3 Contexts

A case $w \in W$ may involve multiple contexts or situations, i.e., different choice-points, different sub-problems, different loci of ethical salience. Let X denote the set of such **contexts**.⁶ A context $x \in X$ is a locus at which agents’ evaluative frames may differ: what vocabulary is locally salient, what calibration is locally operative, may shift from one context to another. The full evaluative landscape is thus parametrized by *both* agents and contexts.

2.2 Local Gauge Transformations

2.2.1 The gauge group

A **gauge group** G is a group (i.e., a collection of transformations that can be composed and reversed) that acts on the space of evaluative frames. Each element $g \in G$ transforms one evaluative frame into another.⁷ What G looks like concretely depends on what kind of frame-variation is at issue. Two examples will serve as running illustrations throughout this section; they differ in both mathematical structure and philosophical character.

Example 1: The scaling group (abelian). Consider the case from §1 where agents calibrate the *magnitude* of ethically relevant features on different numerical scales. One agent might register moderate suffering as the number 3, another as 30, not because they disagree about how much suffering is present, but because their internal “units” differ, much as Celsius and Fahrenheit assign different numbers to the same temperature.

The relevant gauge group here is the **positive real numbers under multiplication**:

$$G = (\mathbb{R}_{>0}, \times).$$

⁶In the continuous physical setting, X corresponds to the spacetime manifold: the set of “points” at which the fields take values. In the discrete ethical setting, contexts are the ethically relevant situations or choice-points within a case. The formalism is indifferent to whether X is finite, countable, or continuous; I work with the finite case throughout.

⁷Recall from the companion paper’s appendix that a *group* is a set G equipped with a composition operation (if $g, h \in G$, then $gh \in G$), an identity element $e \in G$ that changes nothing ($eg = ge = g$), and inverses ($g^{-1}g = gg^{-1} = e$). The composition is associative: $(gh)k = g(hk)$. A group *action* of G on a set S assigns to each $g \in G$ and $s \in S$ a transformed element $g \cdot s \in S$, subject to $e \cdot s = s$ and $(gh) \cdot s = g \cdot (h \cdot s)$.

An element $g \in G$ is simply a positive number: applying it to a frame *rescales* every magnitude by that factor. If $g = 10$, then agent i 's suffering-magnitude of 3 becomes $10 \times 3 = 30$.

This group is **abelian**: the order in which you compose rescalings does not matter ($g_1 \cdot g_2 = g_2 \cdot g_1$ for all $g_1, g_2 \in G$). Rescaling by 2 and then by 5 is the same as rescaling by 5 and then by 2 (both give rescaling by 10). This commutativity will simplify the curvature story considerably, as will become clear in §2.4.

Example 2: Vocabulary transformations (nonabelian). A richer and arguably more paradigmatic case involves agents embedded in different cultural-ethical traditions who organize ethical evaluation through different **conceptual vocabularies**. Consider three traditions:

- A *Confucian* tradition that organizes ethical evaluation around role-constituted relational obligations, e.g., *rén* (benevolence/humaneness), *yì* (righteousness), *lǐ* (ritual propriety), so that the primary ethical categories are structured by the agent's position within a web of relationships.
- A *Kantian* tradition organized around duty, the categorical imperative, and respect for rational agency, where the primary ethical categories concern the form of the maxim (universalizability, treatment of persons as ends).
- An *Aristotelian* tradition organized around character, virtue, and *eudaimonia*, where the primary categories concern the agent's dispositions and their relation to human flourishing.

A gauge transformation here is a systematic **translation map** between vocabularies: it takes an ethical description articulated in one tradition's categories and re-expresses it in another's. Such translations are not arbitrary relabelings; they are structured mappings that preserve certain relationships (e.g., that what one tradition calls an obligation of *rén* maps to something the Kantian tradition recognizes as duty, though the mapping may involve nontrivial restructuring of categories).

The collection of all such vocabulary transformations forms a group G , but (and this is the crucial structural point) it is **nonabelian**: the order in which you compose translations *matters*. Consider the composition:

$$\text{Confucian} \xrightarrow{g_1} \text{Kantian} \xrightarrow{g_2} \text{Aristotelian}$$

versus the direct translation:

$$\text{Confucian} \xrightarrow{g_3} \text{Aristotelian}.$$

The intermediate passage through Kantian categories may systematically reshape certain distinctions. For instance, the relational texture of *rén* may be flattened into the more abstract Kantian category of duty before being re-expanded into the Aristotelian virtue of friendship, losing nuances that a direct

Confucian \rightarrow Aristotelian translation would preserve. The composition $g_2 \circ g_1$ and the direct map g_3 may disagree precisely because the intermediate Kantian framing distorts distinctions that are visible from both the Confucian and Aristotelian standpoints. In formal terms, vocabulary translations do not generally commute: $g_2 \circ g_1 \neq g_1 \circ g_2$.⁸

2.2.2 Local gauge transformations defined

We can now define the central object. A **local gauge transformation** is a function

$$g_x : I \longrightarrow G$$

that assigns to each agent $i \in I$, at each context $x \in X$, a frame-transformation $g_x(i) \in G$. The subscript x indicates that the transformation may vary across contexts; the argument i indicates that it may vary across agents. The defining feature of a *local* transformation is that $g_x(i)$ and $g_x(j)$ may differ for different agents i and j , and $g_x(i)$ and $g_{x'}(i)$ may differ for different contexts x and x' .

Under a local gauge transformation, agent i 's frame at context x shifts from $f_{i,x}$ to $g_x(i) \cdot f_{i,x}$.

Compare this with a **global** gauge transformation, which is the special case where $g_x(i) = g$ for all $i \in I$ and all $x \in X$ —the *same* transformation is applied to every agent at every context. Global invariance, the subject of the companion paper, requires that evaluation be insensitive to such uniform shifts. Local invariance requires insensitivity to the far larger class of transformations where each agent's frame, at each context, may be shifted independently.

2.2.3 The ethical local invariance requirement

We are now in a position to state the ethical local invariance requirement precisely.

Definition 1 (Ethical local gauge invariance) *An evaluation $E : W \rightarrow V$ is **locally gauge-invariant** if and only if, for every local gauge transformation $g_x : I \rightarrow G$, the evaluation rendered after applying g_x to each agent's frame at each context is identical to the evaluation rendered before:*

$$E(g_x(i) \cdot w) = E(w) \quad \forall g_x(i) \in G, w \in W.$$

This states that evaluation depends only on the ethically relevant features themselves, not on the local representational parameters through which any particular agent, at any particular context, registers those features.

⁸In a more detailed formalization, one could model the vocabulary group as a subgroup of the general linear group $GL(n, \mathbb{R})$ acting on a vector space of ethical-category coordinates. Noncommutativity of matrix multiplication then encodes the non-commutativity of vocabulary translations. For present purposes, the essential point is structural: the vocabulary group is nonabelian, and this has consequences for the gauge-invariant content (curvature), as I develop in §2.4.

This requirement is strictly stronger than global invariance. A global gauge transformation is the special case where $g_x(i)$ is the same for all i and x . Local invariance demands invariance under the much larger class of transformations in which the frame-shift can vary from agent to agent and from context to context.

A concrete illustration. Return to the scaling example. Three agents i , j , and k evaluate a situation involving suffering at context x . Each registers the magnitude of the suffering, but on potentially different numerical scales. Agent i might represent moderate suffering as the value 3, agent j as 7, and agent k as 5, not because they disagree about how much suffering is present, but because they calibrate their scales differently. Here $G = (\mathbb{R}_{>0}, \times)$: a local gauge transformation at agent i and context x multiplies i 's numerical assignments by a positive constant $g_x(i)$. Ethical local invariance requires that the ethical evaluation of the situation (who is wronged, what obligations obtain, how agents' claims compare) be insensitive to these local rescalings.

2.2.4 Why local invariance? Two complementary motivations

Before turning to the formal consequences of local invariance, it is worth pausing to make explicit *why* the requirement deserves to be taken seriously. Two complementary considerations converge on the same formal demand.

The cross-comparison motivation. The most immediate motivation is practical: without local invariance, one cannot formulate *cross-agent ethical claims* in a way that satisfies OI. Ethical evaluation is not merely a matter of assessing each agent's situation in isolation; it requires *comparing* agents' claims, weighing their interests, and determining the structure of their mutual obligations. But as §1 observed, any such comparison presupposes that the agents' evaluative descriptions can be brought into a common frame, and if agents' frames vary locally, the comparison is frame-dependent unless local invariance is imposed. The argument of §2.3 below will make this precise: the very attempt to compare evaluative content across differently framed agents forces the introduction of new structure (a connection) that local invariance then constrains.

The action-guiding motivation. A deeper motivation concerns the distinction between *static verdicts* and *action-guiding content*. Global invariance ensures that the *pointwise* evaluative verdicts assigned to cases are objective: the evaluation $E(w)$ does not depend on who the agents are or which global frame is used. But ethics is not merely a matter of assigning static verdicts to cases. It is *action-guiding*: it concerns which transitions between states of affairs are rationally required, permitted, or forbidden, which deliberative moves are admissible, which trade-offs are justified, which paths through the ethical landscape are navigable.

As noted in a **Remark** in the companion paper's formal appendix, this action-guiding content can be thought of as a **directed transition structure** on cases: $w \rightsquigarrow w'$ expresses that the evaluation rationally guides or constrains

a transition from w to w' in deliberation. Two evaluations E_1, E_2 might agree on every pointwise verdict, i.e., $E_1(w) = E_2(w)$ for all w , yet disagree on which transitions they license, because they organize the evaluative landscape differently. The action-guiding content is richer than the pointwise content.

Local invariance, on this deeper understanding, requires invariance not merely of pointwise verdicts but of the full transition structure under locally varying recalibrations. This is precisely why local symmetry in physics constrains *dynamics* (the equations of motion, the laws governing change) rather than merely static descriptions: the gauge field ensures that the dynamical content of the theory is frame-independent, not just the values of individual quantities at individual points.⁹

The two motivations are complementary. The cross-comparison motivation identifies a practical necessity (one needs to compare agents' situations) and shows that local invariance is the condition under which such comparisons are objective. The action-guiding motivation identifies a deeper structural point: that what makes ethics ethically interesting (its capacity to guide action, not merely to assign verdicts) is precisely the kind of content that local invariance protects.

2.3 Derivation of the Ethical Connection Field

The present section shows that the ethical local gauge invariance cannot be achieved “for free.” It requires the introduction of new structure into the evaluative framework (an ethical connection field) that is not present in the global case and that plays a specific, formally characterizable role.

2.3.1 The problem of cross-agent comparison

The difficulty becomes visible as soon as one attempts the most basic operation in ethical evaluation: *comparing* the evaluative situations of two agents.

Consider agents i and j . In the global case (where all agents share a common evaluative frame) comparing their situations is straightforward: both are described in the same frame, and the comparison inherits whatever objectivity the shared frame provides. But once one allows local frame-variation, this straightforward comparison breaks down.

Suppose we want to determine whether agent i 's claim (described in i 's frame) is stronger than, weaker than, or equivalent to agent j 's claim (described in j 's frame). To make this comparison, one must somehow bring the two descriptions into a common frame. But any procedure for doing so involves a *translation* from i 's frame to j 's, and this translation is precisely the kind of structure that a local gauge transformation can alter.

The scaling example. Agent i registers suffering at magnitude 3, agent j at magnitude 7. Is j 's suffering greater? That depends entirely on the relationship

⁹In the companion paper's appendix, this point is developed in a remark on the “directed transition structure” induced by evaluation. The present paper can be read as taking that remark as its point of departure and developing its formal consequences.

between their scales, a relationship that a local gauge transformation can alter. If one rescales i 's frame by a factor of 4 (so i now registers the same suffering as $4 \times 3 = 12$) while leaving j 's frame untouched, the naïve comparison reverses. The comparison, as stated, is *not* locally gauge-invariant.

The vocabulary example. Agent i , reasoning within a Confucian vocabulary, describes an obligation in terms of *rén*. Agent j , reasoning within a Kantian vocabulary, describes an obligation in terms of duty. Are they recognizing the same obligation? Answering this requires a translation between vocabularies, and as noted in §2.2.1, such translations compose noncommutatively. The “same obligation” claim is not well-defined without a specification of *how* to translate, and that specification can shift under a local gauge transformation of either agent’s vocabulary.

This is not a mere technical inconvenience. It is a direct manifestation of the problem §1 identified: without a principled way to compare evaluative content across locally varying frames, one cannot formulate cross-agent ethical claims—claims about the relative weight of different agents’ interests, the comparative urgency of their needs, the structure of their mutual obligations—in a way that satisfies OI. And such cross-agent claims are the very substance of ethics.

2.3.2 The connection field as compensating structure

The solution to the above problem of cross-agent comparison is to introduce a structure that specifies, for each pair of agents (at each context), how to *transport* evaluative content from one agent’s frame to another’s in a way that compensates for local frame-variation.

Definition 2 (Ethical Connection Field) *An ethical connection field on I (at a context $x \in X$) is an assignment \mathcal{C} that associates to each ordered pair of agents (i, j) a group element*

$$\mathcal{C}_{ij} \in G,$$

*interpreted as the frame-adjustment needed to translate evaluative content described in i 's frame into j 's frame.*¹⁰

Think of \mathcal{C}_{ij} as a “conversion factor” or “translation manual” from i 's evaluative language to j 's. In the scaling example, if $\mathcal{C}_{ij} = 7/3$, this means that to translate from i 's numerical scale to j 's, we multiply by $7/3$. In the vocabulary example, \mathcal{C}_{ij} is a systematic mapping from i 's conceptual categories to j 's.

The ethical connection field (also ‘connection field’ or ‘connection’ for short) enters evaluation as follows. When comparing the evaluative situations of agents

¹⁰In the continuous case (standard gauge theory), the connection is a Lie-algebra-valued 1-form on the base manifold, and the transport elements \mathcal{C}_{ij} arise as path-ordered exponentials (*holonomies*) of this form along paths from i to j . In the discrete case, the connection is the assignment of group elements to edges (ordered pairs of agents), and no integration is required. The conceptual role is identical: the connection specifies how to relate local frames at different points of the domain.

i and j , one does not directly compare the raw descriptions in their respective frames. Instead, one first *transports* i 's description to j 's frame by applying \mathcal{C}_{ij} , and then compare the transported description with j 's own description.

Example (scaling). Agent i registers suffering at magnitude 3. Agent j registers suffering at magnitude 7. If $\mathcal{C}_{ij} = 7/3$, then transporting i 's description to j 's frame yields $3 \times (7/3) = 7$ —equal to j 's own registration. The comparison now reveals what we wanted: the underlying suffering is the same despite the different numerical representations.

Example (vocabulary). Agent i describes a situation using the Confucian category of *rén*. The connection \mathcal{C}_{ij} maps this to the Kantian category of “duty arising from respect for persons.” Agent j 's own Kantian description of the same situation also invokes this category. The connection has enabled a cross-vocabulary comparison that would otherwise be ill-defined.

2.3.3 The transformation law for the connection field

The essential question is: *how must the connection field transform under a local gauge transformation if the evaluation is to remain invariant?*

This is the central derivation. Suppose a local gauge transformation $g_x : I \rightarrow G$ is applied. Agent i 's frame is shifted by $g_x(i)$; agent j 's frame is shifted by $g_x(j)$. Before the transformation, the connection element \mathcal{C}_{ij} accomplished the translation from i 's frame to j 's. After the transformation, the translation must accomplish the same work, but now from i 's *shifted* frame to j 's *shifted* frame.

We can read off what the new connection must be by tracking the chain of operations. Starting in i 's shifted frame:

1. *Undo* i 's gauge shift: apply $g_x(i)^{-1}$. This returns us to i 's original frame.
2. *Transport* using the original connection: apply \mathcal{C}_{ij} . This takes us from i 's original frame to j 's original frame.
3. *Apply* j 's gauge shift: apply $g_x(j)$. This takes us from j 's original frame to j 's shifted frame.

Composing these three operations, the transformed connection is:

$$\boxed{\mathcal{C}_{ij} \longmapsto g_x(j) \cdot \mathcal{C}_{ij} \cdot g_x(i)^{-1}}$$

The logic of this transformation law is transparent: it compensates for the local mismatch between $g_x(i)$ and $g_x(j)$. If the transformation were *global* (i.e., $g_x(i) = g_x(j) = g$ for all i, j), then the connection would transform as $\mathcal{C}_{ij} \mapsto g \cdot \mathcal{C}_{ij} \cdot g^{-1}$, which is conjugation by g . In the abelian (scaling) case, where $g \cdot h = h \cdot g$ for all $g, h \in G$, conjugation is trivial and a global transformation leaves the connection unchanged. But for *local* transformations, where $g_x(i) \neq g_x(j)$,

the connection must genuinely adjust to maintain coherent translation between differently shifted frames.¹¹

2.3.4 Three critical observations

Three consequences of this derivation deserve emphasis.

First: the connection field is forced. The connection is not an optional addition to the framework. Without it, cross-agent comparisons are not locally gauge-invariant, and the evaluative framework cannot satisfy the local invariance requirement of §2.2.3. The connection is *necessitated* by the demand for local invariance, just as the gauge field in physics is necessitated by the demand that dynamics be invariant under locally varying phase transformations.

This is the sense in which the local invariance requirement generates content, as anticipated in §1.3. The requirement does not merely constrain the form of existing evaluative operations; it forces the existence of a new structural element—the connection—with determinate properties. Anyone who accepts that objective evaluation must be locally gauge-invariant is thereby committed to the existence of a structure that translates evaluative content across locally varying frames.

Second: the connection field is what makes the directed transition structure locally invariant. The action-guiding motivation of §2.2.4 observed that ethics is not merely a matter of assigning static verdicts to cases but concerns which transitions between evaluative situations are rationally mandated, permitted, or forbidden, a directed transition structure $w \rightsquigarrow w'$ on the space of cases. Local invariance, on that deeper understanding, requires invariance not merely of pointwise verdicts but of this full transition structure under locally varying recalibrations. The connection field is what makes this possible. Without it, the transition structure as described by agent i (in i 's local frame) cannot be coherently compared with the transition structure as described by agent j (in j 's frame), and the question of whether the two agents recognize the same transitions as rationally mandated is not well-posed. The ethical connection field's translation of evaluative content across frames (§2.3.2) is thus not merely a device for comparing static evaluative magnitudes (such as the numerical weight assigned to suffering); it is, more fundamentally, a device for comparing the **dynamical** content of agents' evaluative situations, i.e., the pattern of rational pressures, obligations, and permissible deliberative moves that each agent's frame articulates. Cross-agent comparison of this dynamical content is the very substance of ethical evaluation, and the connection is what renders it frame-independent.

Third: the connection field is not gauge-invariant. The connection's components \mathcal{C}_{ij} shift under local gauge transformations: $\mathcal{C}_{ij} \longmapsto g_x(j) \cdot \mathcal{C}_{ij} \cdot$

¹¹This is the discrete analog of the familiar gauge transformation law for the connection 1-form in continuous gauge theory: $A_\mu \mapsto gA_\mu g^{-1} + g\partial_\mu g^{-1}$. In the discrete case, the “derivative” term is absorbed into the finite group elements, and the transformation law takes the cleaner multiplicative form above.

$g_x(i)^{-1}$. This means that the connection cannot, by itself, constitute the objective ethical content we are after. It is partly representational, partly substantive: it depends in part on the (arbitrary) local frames, and in part on something deeper. The question then becomes: *what part of the connection's structure is gauge-invariant?* What content does the connection carry that survives every local re-framing? That question leads us to curvature.

2.4 Derivation of the Ethical Curvature/Field Strength: The Gauge-Invariant Content

The answer to the question just posed is *curvature* or *field strength*: the gauge-invariant content of the connection, which measures the extent to which the connection's comparison procedure is *path-dependent*.¹²

2.4.1 Path-dependence and the holonomy around a loop

Consider three agents i , j , and k (at some context x). There are two natural ways to transport evaluative content from i 's frame to k 's:

- **Directly:** apply \mathcal{C}_{ik} , transporting from i to k in one step.
- **Indirectly:** first transport from i to j (using \mathcal{C}_{ij}), then from j to k (using \mathcal{C}_{jk}). The indirect transport is the composition $\mathcal{C}_{jk} \cdot \mathcal{C}_{ij}$.

If these two procedures always agree (i.e., $\mathcal{C}_{ik} = \mathcal{C}_{jk} \cdot \mathcal{C}_{ij}$ for every triple of agents) then the connection is **flat**: cross-agent comparison is independent of the path of intermediate comparisons used to establish it.

A flat connection is, in a precise sense, *eliminable*. It can be generated by a choice of “absolute” frame for each agent: a single reference frame such that each agent's local frame is a definite transformation of the reference. All the comparative information the connection encodes reduces to the relationships between agents' frames and this shared reference. In the scaling example, a flat connection means there exists a single “true” scale such that each agent's local scale is a definite multiple of it; the connection merely records these multiples.

But a flat connection is a very special case. In general, the two transport procedures *differ*, and the discrepancy measures the **holonomy** around the

¹²The terms “curvature” and “field strength” refer to the same mathematical object. In the physics literature, “curvature” foregrounds the geometric interpretation (path-dependence of parallel transport), while “field strength” foregrounds the dynamical interpretation (the quantity that governs how the field acts on matter, mediates forces, and determines which physical transitions occur). For the ethical application, both interpretations are relevant: the curvature registers the irreducible path-dependence of cross-agent comparison, while the field strength registers the *structure of rational pressure across cases*, i.e., which evaluative transitions are mandated, which trade-offs are admissible, how evaluative demands change as one moves through the ethical landscape. Because this dynamical dimension is central to the ethical significance of the gauge-invariant content, I use both terms throughout, with “field strength” when emphasizing the action-guiding, dynamical character of the content, and “curvature” when emphasizing its geometric structure.

loop $i \rightarrow j \rightarrow k \rightarrow i$: the net frame-transformation accumulated by transporting evaluative content around a closed circuit of agents and returning to the starting point.

Definition 3 (Ethical Curvature/Field Strength) *The ethical curvature (or field strength) of the connection field \mathcal{C} at the loop (i, j, k) is:*

$$\mathcal{E}_{ijk} = \mathcal{C}_{ki} \cdot \mathcal{C}_{jk} \cdot \mathcal{C}_{ij}$$

This is the ordered product of connection elements around the closed loop $i \rightarrow j \rightarrow k \rightarrow i$.¹³ If $\mathcal{E}_{ijk} = e$ (the identity element of G , which represents “no transformation”), then transport around this loop is trivial: evaluative content returns to its original form after the circuit. If $\mathcal{E}_{ijk} \neq e$, the connection has nontrivial structure at this loop—structure that cannot be eliminated by any local choice of frames.

Example (scaling). Agents i, j, k with $G = (\mathbb{R}_{>0}, \times)$. Suppose $\mathcal{C}_{ij} = 2$ (to go from i ’s scale to j ’s, multiply by 2), $\mathcal{C}_{jk} = 3$ (from j to k , multiply by 3), and $\mathcal{C}_{ki} = 1/5$ (from k back to i , multiply by $1/5$). Then

$$\mathcal{E}_{ijk} = \mathcal{C}_{ki} \cdot \mathcal{C}_{jk} \cdot \mathcal{C}_{ij} = \frac{1}{5} \times 3 \times 2 = \frac{6}{5} \neq 1.$$

The curvature is nontrivial: the scales do not “close up” coherently around the loop. If the connection were flat (all scale-relationships globally coherent), we would have $\mathcal{E}_{ijk} = 1$. The departure from 1 measures the irreducible discrepancy—the extent to which the agents’ scale-relationships resist being generated by a single reference scale.

Example (vocabulary). Agents i (Confucian), j (Kantian), k (Aristotelian). The connection elements $\mathcal{C}_{ij}, \mathcal{C}_{jk}, \mathcal{C}_{ki}$ are the vocabulary translation maps. The curvature $\mathcal{E}_{ijk} = \mathcal{C}_{ki} \cdot \mathcal{C}_{jk} \cdot \mathcal{C}_{ij}$ measures the net “distortion” accumulated by translating an ethical description from the Confucian vocabulary to the Kantian, then to the Aristotelian, then back to the Confucian. If $\mathcal{E}_{ijk} = e$ (the identity translation), the circuit is lossless: the three traditions’ vocabularies are globally coherent, and all cross-traditional differences are eliminable by choosing a master vocabulary. If $\mathcal{E}_{ijk} \neq e$, there is an irreducible discrepancy, i.e., a structural feature of how the three traditions’ conceptual schemes interrelate that cannot be dissolved by any single re-framing.

2.4.2 The gauge-invariance of curvature/field strength

The crucial property of \mathcal{E}_{ijk} is its behavior under local gauge transformations. Using the transformation law for the connection: $\mathcal{C}_{ij} \mapsto g_x(j) \cdot \mathcal{C}_{ij} \cdot g_x(i)^{-1}$,

¹³In continuous gauge theory, the holonomy is the path-ordered exponential of the connection 1-form around a closed loop, and curvature (the field strength tensor) is the infinitesimal holonomy around an infinitesimal loop. In the discrete case, holonomy around a loop of agents is the ordered product of connection elements around the loop, and curvature is the holonomy (or equivalently, its deviation from the identity element $e \in G$). The conceptual content is the same: curvature measures the irreducible path-dependence of parallel transport.

one can verify by direct computation that:¹⁴

$$\boxed{\mathcal{E}_{ijk} \mapsto g_x(i) \cdot \mathcal{E}_{ijk} \cdot g_x(i)^{-1}}$$

The curvature transforms by **conjugation** at the base point i of the loop. This has two immediate consequences, which differ in an illuminating way between the abelian and nonabelian cases.

Abelian case (e.g., the scaling group $G = (\mathbb{R}_{>0}, \times)$). When G is abelian, conjugation is trivial: $g \cdot h \cdot g^{-1} = h$ for all $g, h \in G$. Therefore:

$$\mathcal{E}_{ijk} \mapsto \mathcal{E}_{ijk}.$$

The curvature/field strength itself is fully gauge-invariant. Under any local rescaling of agents' frames, the curvature does not change at all. This makes the scaling example especially transparent: the number $\mathcal{E}_{ijk} = 6/5$ in the example above is an *absolute*, frame-independent fact about the relationship among the three agents' scales. No local rescaling can alter it.

Nonabelian case (e.g., vocabulary transformations). When G is nonabelian, the curvature transforms by conjugation, which is nontrivial: the “matrix representation” of \mathcal{E}_{ijk} depends on the choice of local frame at the base point i . However, the **conjugacy class** of \mathcal{E}_{ijk} —and in particular, whether \mathcal{E}_{ijk} is trivial or nontrivial, and all its algebraic invariants (traces, characteristic polynomials, eigenvalue spectra)—*does not depend on any local frame choice*.¹⁵

This is the precise sense in which the nonabelian case is richer. The curvature is not fully gauge-invariant (its representation depends on the local frame), but its *invariant content* (the content extractable via conjugation-invariant functions) is entirely gauge-independent. In the vocabulary example, the specific “matrix” describing the net distortion around the Confucian \rightarrow Kantian \rightarrow Aristotelian circuit depends on which tradition's categories we use to describe it (i.e., the choice of base-point frame). But *whether* there is net distortion, and

¹⁴The computation is instructive. Under a local gauge transformation, $\mathcal{C}_{ij} \mapsto g_x(j) \cdot \mathcal{C}_{ij} \cdot g_x(i)^{-1}$. Substituting into $\mathcal{E}_{ijk} = \mathcal{C}_{ki} \cdot \mathcal{C}_{jk} \cdot \mathcal{C}_{ij}$:

$$\begin{aligned} \mathcal{E}_{ijk} &\mapsto [g_x(i) \cdot \mathcal{C}_{ki} \cdot g_x(k)^{-1}] \cdot [g_x(k) \cdot \mathcal{C}_{jk} \cdot g_x(j)^{-1}] \cdot [g_x(j) \cdot \mathcal{C}_{ij} \cdot g_x(i)^{-1}] \\ &= g_x(i) \cdot \mathcal{C}_{ki} \cdot \underbrace{g_x(k)^{-1} \cdot g_x(k)}_{=e} \cdot \mathcal{C}_{jk} \cdot \underbrace{g_x(j)^{-1} \cdot g_x(j)}_{=e} \cdot \mathcal{C}_{ij} \cdot g_x(i)^{-1} \\ &= g_x(i) \cdot \mathcal{E}_{ijk} \cdot g_x(i)^{-1}. \end{aligned}$$

The interior gauge factors cancel in a telescoping pattern, leaving only conjugation by $g_x(i)$ at the base point of the loop.

¹⁵In the mathematical theory of groups, two elements h and h' of a group G are *conjugate* if there exists $g \in G$ such that $h' = ghg^{-1}$. Conjugacy is an equivalence relation on G , and the equivalence classes are called *conjugacy classes*. A function on G is *conjugation-invariant* (or a *class function*) if it takes the same value on all elements of a given conjugacy class. Examples include the trace and determinant of a matrix representation. These conjugation-invariant quantities are the gauge-invariant objective magnitudes of the theory.

how much (quantified by invariant measures like traces), does not depend on this choice.

Directed transition structure on cases. The gauge-invariance of the curvature has a direct consequence for the directed transition structure $w \rightsquigarrow w'$ on the space of cases (introduced in §2.2.4). As the second observation of §2.3.4 noted, the connection makes it possible to *compare* the transition structures that different agents' frames articulate, to ask, in a frame-independent way, whether agents recognize the same evaluative transitions as rationally mandated. The curvature/field strength answers a further question: what is the **invariant content** of this transition structure?

Consider the loop $i \rightarrow j \rightarrow k \rightarrow i$. Each leg of the loop involves the connection translating an evaluative description (including the rational pressures and permissible transitions it encodes) from one agent's frame to the next. If the curvature/field strength \mathcal{E}_{ijk} is trivial, the transition structure is globally coherent: what counts as a rationally mandated transition from i 's perspective, once translated through j and k and back, returns unchanged. There is, as it were, a single pattern of rational pressure that all agents articulate in their local idioms. If the curvature/field strength is nontrivial, there is an irreducible discrepancy in the transition structure itself, a feature of how the agents' patterns of rational pressure interrelate that no local re-framing can dissolve. The curvature/field strength thus captures not merely the path-dependence of static cross-agent comparison but the gauge-invariant dynamical content of the ethical landscape: the objective structure of which evaluative transitions are mandated, which trade-offs are admissible, and how rational pressure changes as one moves through the space of cases. It is this dynamical content (not any particular agent's local description of it) that constitutes the ethically real transition structure.

2.4.3 Summary: what the curvature/field strength tells us

Let us take stock. The connection components \mathcal{C}_{ij} are gauge-dependent, i.e., they shift under local re-framings. The curvature invariants of the connection are gauge-independent, i.e., they survive every local re-framing. This structural fact carries a direct philosophical lesson:

If there is genuine ethical content encoded in the way agents' evaluative situations are related, and if that content must be objective in OI's sense, then it lives in the curvature or field strength \mathcal{E}_{ijk} , not in the connection components \mathcal{C}_{ij} . The field strength encodes the structure of rational pressure across cases—how evaluative demands change, which transitions are mandated, which trade-offs are admissible—independently of any local representation. This dynamical structure is the ethically real content; not any particular evaluative assignment in a given frame.

2.4.4 Philosophical interpretation of the curvature

Consider what curvature *means* in the ethical context. The flat and curved cases carry sharply different philosophical implications, and it is worth developing both.

2.4.4.1 Flat connections and the “ethical Esperanto” hypothesis

A **flat connection** among agents i , j , and k would mean that the comparison relations among their evaluative frames are globally coherent: if we know how i 's frame relates to j 's, and how j 's relates to k 's, we can deduce how i 's relates to k 's with no residual discrepancy. This is the case where all frame-variation is, as it were, merely perspectival, i.e., eliminable by a single global re-calibration.

In the scaling case, flatness means there exists a “master scale”: a single canonical numerical calibration from which each agent's local scale is derivable by a definite multiplicative factor. In the vocabulary case, flatness means something more philosophically striking: there exists what we might call an “**ethical Esperanto**”: a single canonical evaluative vocabulary from which every tradition's conceptual categories can be derived by a definite (and globally coherent) translation. The Confucian vocabulary, the Kantian vocabulary, and the Aristotelian vocabulary would then be related to this master vocabulary exactly as Celsius, Fahrenheit, and Kelvin are related to the “true” temperature: different representational schemes for the same underlying content, with fully coherent conversion rules.

Flatness is thus a *strong* and philosophically non-trivial condition. It amounts to the claim that the ethical terrain, however complex, admits of a single unified description from which all local descriptions can be recovered. This is a version of what one might call *evaluative monism*: the thesis that there is, in principle, a single correct way of articulating the ethically relevant features of any situation, and that all other articulations are perspectival re-descriptions of it.

Whether the connection of the actual ethical community is flat is an open question and a deeply consequential one. But the framework does not presuppose either answer. It provides a formal criterion for distinguishing the flat case from the curved case, and it identifies what is at stake in each.

2.4.4.2 Nontrivial curvature: irreducible relationality and irreducible complexity

Nontrivial curvature means that no ethical Esperanto exists. The relations among agents' evaluative frames have an irreducible structure, a structure that cannot be captured by any assignment of “absolute” evaluative positions to individual agents, and that cannot be dissolved by any choice of master vocabulary or master scale.

This carries two distinct philosophical lessons.

Irreducible relationality. Nontrivial curvature is a substantive claim about

the ethical community, not about individual agents. It says that the network of evaluative relationships among agents has *objective, gauge-invariant content* that is not reducible to facts about individual agents taken separately. The curvature registers a *holistic* property of ethical life: how agents’ evaluative situations interrelate in a way that no local re-framing can dissolve.

More concretely, nontrivial curvature registers the presence of genuine ethical structure that requires attending to the *pattern* of cross-agent relationships, not merely to each agent’s situation in isolation. It is the formal counterpart of a familiar insight from lived ethical practice: ethical life is irreducibly relational. The ethical landscape is not a collection of independent evaluative positions but a structured web of comparisons, obligations, and claims whose content cannot be decomposed into agent-by-agent facts without remainder.

Irreducible complexity. But nontrivial curvature signals something further: an irreducible *complexity* in the ethical terrain. If the connection were flat, then despite the surface diversity of agents’ evaluative frames, the underlying ethical structure would be simple in a precise sense: fully encodable in a single vocabulary, fully decomposable into agent-by-agent facts expressed in that vocabulary. Nontrivial curvature means the structure resists this compression. The ethical terrain has a kind of complexity that is not an artifact of how we describe it (since the curvature invariants are frame-independent) but is intrinsic to the ethical relationships themselves.

This is not complexity in the sense of mere *difficulty*—the practical difficulty of figuring out what to do in hard cases. It is structural complexity: the ethical community’s evaluative relationships have an internal geometry that cannot be flattened out. In the vocabulary case, this means that the web of inter-traditional relationships carries information that no single tradition can fully articulate on its own—each tradition “sees” the ethical landscape from a local vantage point, and the curvature measures what is visible only from the vantage point of the relationships *among* traditions, not from within any single one.¹⁶

In the nonabelian case, this insight acquires additional texture. The curvature transforms by conjugation ($\mathcal{E}_{ijk} \mapsto g_x(i) \cdot \mathcal{E}_{ijk} \cdot g_x(i)^{-1}$), so its invariant content (extracted via conjugation-invariant functions) is entirely gauge-independent. The specific “matrix” describing the net distortion around the Confucian \rightarrow Kantian \rightarrow Aristotelian circuit depends on which tradition’s categories we use to describe it (the choice of base-point frame). But *whether* there is net distortion, and *how much* (quantified by invariant measures), does not depend on this choice. The irreducible complexity is objectively there, no matter where you stand.

Dynamical counterparts. The two lessons just identified (irreducible relationality and irreducible complexity) have been stated primarily in the geometric

¹⁶This gives formal expression to a thought that comparative ethicists have long found compelling: that the relationships *among* ethical traditions are themselves ethically significant, not merely instrumental to the project of identifying the “correct” tradition from which to evaluate. The curvature framework makes this intuition precise and identifies its gauge-invariant content.

register: they concern the structure of cross-agent **comparison**. But each has a dynamical counterpart that concerns the structure of rational **pressure**, and the directed transition structure introduced in §2.2.4 makes these counterparts visible. Irreducible relationality, read dynamically, means that the pattern of rational pressure across agents (which transitions are mandated, permitted, and forbidden for whom, which trade-offs are admissible between whose claims) is not decomposable into agent-by-agent rational pressures determined in isolation, i.e., independently of how other agents' evaluative situations bear on them. The rational pressure on agent i is constitutively shaped by the structure of cross-agent relations (mediated by the connection), and the curvature registers the extent to which this shaping is irreducible. Irreducible complexity, read dynamically, means that the action-guiding content of the ethical landscape, i.e., the full pattern of which evaluative transitions are permitted, mandated, and forbidden and how rational pressure changes as one moves through the space of cases, resists compression into any single agent's local articulation of that content. No agent's frame, however rich, can fully encode the dynamical structure that the curvature captures; the objective action-guiding content of the moral community's evaluative life is, in a precise sense, richer than any individual agent's description of it.

The task of §3 is to show that the structure identified here, i.e., the connection that enforces coherence across locally varying frames, whose gauge-invariant curvature constitutes the objective ethical content, can be identified with what Kant calls the *good will*, and that the formal properties derived in this section illuminate why Kant ascribes to it unconditional value.

3. The Unconditional Shining of the Good Will

The formal apparatus of §2 derived three structural results from the ethical local gauge invariance requirement:

1. Ethical local invariance *necessitates* an ethical connection field \mathcal{C} : a structure that specifies how to transport evaluative content across agents' locally varying frames (§2.3).
2. The ethical connection field itself is not gauge-invariant: its components \mathcal{C}_{ij} shift under local re-framings (§2.3.4).
3. The gauge-invariant content of the ethical connection field resides in its curvature or field strength \mathcal{E} : the irreducible path-dependence of cross-agent comparison/structure of rational pressure across cases (§2.4).

The present section argues that this formally derived structure can be identified with what Kant calls the *good will*, and that the formal properties established in §2 illuminate (with a precision not previously available) why Kant ascribes to it unconditional value.

The argument proceeds in four steps. First, I characterize the structural role that the ethical connection field plays and ask what kind of ethical structure could fill that role (§3.1). Second, I argue that the good will, understood as the orientation of practical reason toward the moral law, fits this structural role (and that the moral law can be understood as an ethical local invariance requirement) (§3.2).¹⁷ Third, I show that the unconditional value Kant attributes to the good will corresponds to gauge-invariance: it is precisely the evaluative content that survives every local re-framing (§3.3). Fourth, I note what this identification illuminates about Kant’s own claims (§3.4).

3.1 The structural profile of the ethical connection field

Before identifying the connection field required by local ethical gauge invariance with anything in the ethical tradition, it is helpful to characterize abstractly what kind of structure it must be. The formal derivation of §2 tells us that the ethical connection field has four defining features:

(a) It enforces coherence across locally varying frames. The connection is what makes it possible for agents with different evaluative calibrations, i.e., different scales, different vocabularies, different ways of registering ethical salience) to be subject to a common ethical evaluation. Without the connection, cross-agent comparison breaks down and local invariance fails. The connection is, in this precise sense, the structure that *holds the moral community together* across its internal representational diversity.

(b) It is necessitated by the demand for objectivity. The connection is not an optional theoretical posit. It is forced by the requirement that objective evaluation be insensitive to locally varying representational parameters. Anyone who accepts local gauge invariance is committed to its existence, whether or not they have a name for it.

(c) It is not itself fully objective. The connection’s components transform under local gauge transformations ($\mathcal{C}_{ij} \mapsto g_x(j) \cdot \mathcal{C}_{ij} \cdot g_x(i)^{-1}$). This means the connection is partly representational: its specific expression depends on

¹⁷A note on terminology is in order. Throughout §2, I refer to the “ethical connection field,” “ethical local gauge invariance,” and “ethical curvature,” using “ethical” as the domain-marker for evaluative structures concerning practical reason. Beginning in §3, I shift to “moral” in certain key phrases (e.g., “moral law,” “moral community,” “moral content”). This shift is deliberate. The usage tracks a distinction in Kant’s own framework. For Kant, the moral law is not simply one law among others that happens to concern the ethical domain; it is the fundamental law of pure practical reason, and its fundamentality and normative authority earn it the title “moral.” Analogously, the connection field derived in §2 operates within the domain of ethical evaluation, but it earns the designation “moral” by virtue of its structural role: it is the structure that makes objective ethical evaluation possible in the first place, not one more ethical structure among others. The curvature it generates is “moral content” in this same sense: it is the gauge-invariant structure whose existence is constitutive of the ethical domain’s objectivity, and it is, for that reason, authoritative. The terminological shift from “ethical” to “moral” thus marks not a change of subject but a recognition of fundamentality: what begins as a structural feature of the ethical domain earns the label ‘moral’ through the fact that it makes ethical objectivity possible and is thus normatively authoritative.

the local frames in which it is described. It does genuine work. Without it, the evaluative framework is incoherent. However, its work is not visible as a frame-independent “object” in the way that the curvature is.

(d) It generates gauge-invariant dynamical content (constituted by its curvature or field strength). The connection’s objective contribution to the ethical landscape is its curvature or field strength: the irreducible, frame-independent structure that emerges from the pattern of cross-agent comparisons it mediates. Crucially, this gauge-invariant content is not merely a static residue of cross-agent comparison. It is *dynamical*: it encodes the structure of rational pressure across cases, i.e., how evaluative demands change as one moves through the ethical landscape, which transitions between evaluative situations are mandated or forbidden, which trade-offs are admissible. This structure of rational pressure is the ethically real content that the connection carries, and it is independent of any particular agent’s local evaluative assignment. The connection is, as it were, the *source* of objective ethical content without itself being fully objective; and the content it sources is not a set of pointwise verdicts but a dynamical structure governing how evaluative life unfolds.

Any ethical structure identified with the connection must exhibit all four features. It must enforce coherence across agents’ diverse evaluative standpoints; it must be a structural requirement of objective practical reason, not an optional moral add-on; its specific expression must vary across agents and contexts without this variation undermining its role; and it must generate genuinely objective ethical content that transcends any particular local framing.

3.2 The good will as the ethical connection field

Kant opens the *Groundwork of the Metaphysics of Morals* with the claim that the good will is the only thing that can be regarded as good “without limitation [*ohne Einschränkung*].”¹⁸ The good will, for Kant, is not a particular desire, emotion, or temperament. It is the *orientation* of practical reason toward acting from duty (i.e., acting only on maxims that one can rationally will as universal law. It is the will insofar as it is determined by the moral law rather than by inclination, self-interest, or the contingencies of the agent’s particular situation.

I now argue that the good will, so understood, satisfies the structural profile of the ethical connection field.

(a) The good will enforces coherence across locally varying frames. Consider what the good will *does* in Kant’s framework. When an agent acts from the good will, they act on a maxim that they can will as a universal law, i.e., a maxim whose prescriptive content does not depend on the agent’s particular position, interests, or evaluative calibration. The universalizability test is precisely a test of *frame-independence*: it asks whether the maxim’s

¹⁸Kant, *Groundwork of the Metaphysics of Morals*, 4:393. All references to Kant are to the standard Akademie edition pagination.

action-guiding content survives the transition from the agent's local evaluative standpoint to the standpoint of any rational agent whatsoever.

This is the work of a connection field. The good will is the structure of practical reason that enables an agent embedded in a particular evaluative frame, with particular projects, relationships, cultural vocabulary, and affective calibration, to generate ethical content that is communicable across frames. It translates the locally framed deliberative situation into content that other agents, with different frames, can recognize as bearing on them. Without the good will (or something playing its structural role), each agent's ethical reasoning remains imprisoned in its local frame, and cross-agent ethical claims lose their objectivity.

(b) The good will is necessitated by the demand for objectivity. The derivation in §2 showed that the connection is not an optional addition to the evaluative framework but is forced by the requirement of local gauge invariance. Analogously, Kant's argument in the *Groundwork* can be read as showing that the good will is not an optional moral ideal but is forced by the structure of practical reason itself, specifically, by the demand that ethical evaluation be objective in the sense of being independent of the agent's contingent inclinations and local evaluative standpoint.¹⁹

The parallel is precise. Just as the gauge connection in physics is forced by the demand that dynamics be invariant under locally varying phase transformations, the good will is forced by the demand that ethical evaluation be invariant under locally varying evaluative re-framings. In both cases, the structural requirement generates the structure: the demand for local invariance does not merely *constrain* an independently given connection but *calls it into existence*.

(c) The good will's expression varies across agents and contexts. A crucial feature of the connection, noted in §2.3.4, is that its components C_{ij} are not gauge-invariant. They shift under local re-framings. This might initially seem like a defect. How can a partly representational structure be identified with something Kant regards as having unconditional value?

But this feature corresponds to a well-known aspect of the good will. The good will does not manifest identically in every agent and every context. Its specific expression (the particular maxims it generates, the particular duties it recognizes, the particular form its moral reasoning takes) varies with the agent's circumstances, capacities, relationships, and (crucially) evaluative vocabulary. A Confucian agent acting from the good will may articulate their duty in terms

¹⁹This reading of Kant is consonant with the "from-grounds" interpretation of the a priori. That is, the good will is not merely known prior to or independently of experience. Rather, it is known *a priori* in the sense that the argument from the *a priori* structure of practical reason (and its constitutive aspiration to objectivity) explains *why* something with the good will's structural properties must exist, if genuinely objective (i.e., categorical) morality is to be possible (cf. Houston Smit's pioneering work on the from-grounds a priori and the condition-constitutive reading of Kant's critical philosophy; see Smit (2009, forthcoming)). The derivation of the connection from local invariance has precisely this character: it shows why a structure with the connection's properties is necessitated by the requirement of objective evaluation, thereby explaining the ground of its possibility.

of *rén*; a Kantian agent may articulate the same duty in terms of the categorical imperative; an Aristotelian agent may articulate it in terms of virtuous character. These are different local expressions of the same underlying orientation of practical reason.

The connection's transformation law ($\mathcal{C}_{ij} \mapsto g_x(j) \cdot \mathcal{C}_{ij} \cdot g_x(i)^{-1}$) formalizes this. Under a shift in evaluative vocabulary (a local gauge transformation), the connection's components change, i.e., the particular "translation manual" between agents' frames is re-expressed in the new vocabulary, but its structural role is preserved. The good will, like the ethical connection field, is not a single fixed object but a structure whose specific expression is frame-dependent while its structural function (enforcing coherence across frames) is not.

(d) The good will generates objective dynamical ethical content. The connection generates gauge-invariant content through its curvature or field strength: the irreducible structure that survives every local re-framing. If the good will is the connection, then the gauge-invariant ethical content it generates—the field strength—is not merely the static residue of cross-agent comparison. It is the *structure of rational pressure across cases*: the objective pattern of which evaluative transitions are mandated, which deliberative moves are admissible, which trade-offs are justified, and how ethical demands change as one moves through the landscape of evaluable situations or cases. This dynamical structure—how evaluative life is governed, not merely how it is described—is the ethically real content, independent of any particular agent's local framing.

Kant's own language is suggestive here. The good will, he says, "shines by its own light, as something that has its full worth in itself" (*Groundwork*, 4:394). It is not made valuable by what it produces (by the consequences of the actions it motivates) but has value in itself, independently of outcomes. I am now in a position to give this claim a precise formal interpretation.

The good will as a single field. The gauge-theoretic identification suggests a further insight that is worth making explicit. In physics, the gauge field is a *single* field defined over the entire base space: it is not a collection of separate fields, one per spacetime point, but a unified structure whose values at different points are related by the connection. Individual charged particles *couple to* and *source* this field—their presence and motion contribute to determining its configuration—but the field is not reducible to any individual particle's contribution. It is a genuinely global structure, constituted by the pattern of local interactions but not decomposable into them.

The ethical analog is striking. If the good will is a connection, then there is, in a precise sense, a *single good will field* defined over the moral community—a single structure of cross-agent evaluative coherence and rational pressure across cases—rather than a collection of separate "good wills," one per agent. Each finite rational being *sources* this field by exercising practical reason to determine their will according to the moral law: each agent's moral activity contributes to constituting the connection, shaping the structure that enforces evaluative coherence across the community. But no individual agent *is* the field. The good will, as connection, is irreducibly communal: it is the structure that relates

agents' locally varying evaluative standpoints, and it exists only insofar as agents sustain it through the ongoing exercise of objective practical reason.

This gives formal expression to a thought latent in Kant's own framework. The moral law is universal (it addresses every rational being as such), and the good will is the will determined by the moral law. If the moral law is (as the present framework suggests) a local invariance requirement on practical reason, then the good will it generates is necessarily a single, community-wide field: the invariance requirement is one requirement, and the connection it necessitates is one connection, sourced by every rational agent who orients their practical reason toward objectivity. The plurality of agents does not generate a plurality of good wills; it generates a plurality of *sources* for a single good will field, just as the plurality of charged particles in physics generates a plurality of sources for a single electromagnetic field.²⁰

3.3 Unconditional value as gauge-invariance

The identification of the good will with the ethical connection field, together with the formal results of §2, yields a precise characterization of what Kant means by the good will's unconditional value.

Consequences are frame-dependent. Consider the consequences of an action: its effects on agents' welfare, its impact on the distribution of goods, its causal ramifications in the world. In the framework of §2, consequences are naturally modeled as features of individual agents' evaluative situations: how much welfare agent *i* gains or loses, described in *i*'s local frame. As such, the description of consequences is *frame-dependent*: it shifts under local gauge transformations. Agent *i*'s welfare-gain of 3, described in *i*'s scale, becomes a welfare-gain of 30 when *i*'s frame is rescaled by a factor of 10. The *numerical magnitude* of the consequence depends on the local calibration.

This does not mean consequences are “unreal” or ethically irrelevant. It means that the ethical significance of consequences, insofar as it is to be objective, must be extracted via gauge-invariant operations, operations that filter out the frame-dependent representation and retain only what is common across all frames. Consequences *contribute* to the objective ethical content, but they do so only when processed through the ethical connection field (which translates them across frames) and evaluated via the curvature/field strength (which extracts the frame-independent content).

The good will's value is gauge-invariant. The good will, as a connection field, is the structure that makes this processing possible. Its value does not reside in any particular consequence it produces (a frame-dependent description) but in its structural role: enforcing coherence across locally varying frames

²⁰This should not be confused with the claim that the good will is a collective entity in the sense of social ontology. The good will field is constituted by individual agents' exercises of practical reason, not by any supra-individual agent. The point is structural: the connection is defined over the space of agents, not within any individual agent, and its existence depends on the pattern of cross-agent relations that individual moral activity sustains.

and thereby generating the gauge-invariant content (curvature) that constitutes objective ethical structure.

The unconditional value of the good will is thus its **gauge-invariant character**: it is precisely the evaluative content that does not depend on any local choice of frame. The good will’s value is unconditional because it is not *conditioned* by the representational parameters of any particular evaluative standpoint. It “shines by its own light” because its light, i.e., its contribution to the ethical landscape is not borrowed from the local frame in which it happens to be expressed. It is visible from every frame, because it is the structure that makes cross-frame visibility possible in the first place. Consequences, outcomes, particular evaluative weights, and other specific substantive ethical considerations are the “ethical matter” fields. The good will as connection is not in competition with them; it’s the structure that makes their coherent comparison across frames possible. Its unconditionality is grounded in its structural role: it’s the condition of frame-independent evaluation, not one more thing to be evaluated.

The jewel analogy, formally. Kant’s celebrated analogy is illuminating: the good will, “like a jewel,” shines “by its own light” even when it “achieves nothing” and “accomplishes nothing” (*Groundwork*, 4:394). In the gauge-theoretic framework, this corresponds to a precise structural fact. The curvature/field strength of the connection (the gauge-invariant content) is a property of the connection *itself*, not of the outcomes it brings about. A connection can have dynamical content independently of whether it couples to any physical system (and thus effects any change). The good will’s “shining” is the dynamical content: the irreducible, frame-independent ethical structure that the connection carries, independently of what happens at any particular agent’s local evaluative position.

3.4 What the formalization illuminates

The gauge-theoretic formalization of the moral law as a local invariance principle and the good will as the connection field required by this principle illuminates several features of Kant’s account that are often regarded as puzzling or under-motivated.

Why the purely formal character of the moral law is a feature. This gauge-theoretic interpretation of the foundations of Kant’s moral philosophy explains why the merely formal or “empty” character of the moral law is not a deficiency but a structural feature. In physics, nobody objects that the local invariance requirement is “empty” because it doesn’t specify the values of the electromagnetic field. The requirement’s generality is precisely its power: it constrains the form of the dynamics without dictating specific field configurations. The specific configurations are determined by initial conditions, boundary conditions, and the coupled dynamics of connection and matter fields. Similarly, the moral law’s generality — its refusal to specify which particular actions are right — is not a failure to provide content but a reflection of its structural

role. It constrains the form of admissible evaluation and generates the connection (good will) that makes locally invariant evaluation possible. The specific moral content (which actions are required in which situations) is determined by the coupled dynamics of the connection and the “ethical matter fields,” i.e., the substantive considerations (welfare, rights, relationships, virtues) that the connection organizes into frame-independent evaluations.²¹

Why the good will is the only unconditionally good thing. Kant’s claim that the good will is the only thing good “without limitation” has struck many readers as dogmatic or question-begging. The gauge-theoretic framework provides a structural rationale. Other candidate goods (talents, temperaments, consequences, even happiness) are naturally modeled as features of agents’ local evaluative situations. As such, they are frame-dependent: their evaluative significance shifts under local re-framings. Only the connection, i.e., the structure that enforces coherence *across* frames, has the gauge-invariant character that “without limitation” requires. The good will is unconditionally good because it is the only structure whose value is constitutively independent of the local frame from which it is assessed.²²

Why the good will is not reducible to good consequences. Consequentialist critics have long pressed the question: if the good will’s value is not its consequences, what *is* its value? The gauge-theoretic answer is precise. Consequences are frame-dependent outcomes at individual agents’ evaluative positions. The good will’s value is the gauge-invariant structure (curvature) that makes it possible to compare, evaluate, and ethically relate those frame-dependent outcomes across agents. The relationship is not rivalry but dependence: consequences matter objectively because there is a connection that renders them cross-agent comparable, and the connection’s own value consists in enabling precisely this comparability. Reducing the good will’s value to consequences would be like reducing the gauge field’s significance to the particular field values at individual spacetime points. It would miss the structural contribution that makes the objectivity of local values and their dynamical relations possible in the first place.

Why reason inevitably recognizes the good will’s value, even in “the most hardened scoundrel.” In the third section of the *Groundwork*, Kant

²¹Although this paper is not the place to develop this point, I argue in other work that the fundamental ethical dimensions of evaluation (axiological, deontological, and aretaic) can fruitfully modeled as dynamically coupled rank-2 tensor fields (which can be represented as two-dimensional matrices in which the diagonal entries represent different ethical dimensions and the off-diagonal entries represent synergies and antagonisms between ethical dimensions) (Sanchez Borboa, manuscript B). Note that this tensor-field model is a further specification of what the ethical matter fields look like, not a correction to the curvature/field strength account developed in this paper.

²²This is not to say that talents, temperaments, and consequences are valueless. They are *conditionally* valuable: their objective ethical significance is mediated by the connection (which translates them across frames) and extracted via the curvature (which identifies the frame-independent content). This is the gauge-theoretic analog of Kant’s own claim that talents and temperaments are good only when governed by a good will.

claims, “There is no one, not even the most hardened scoundrel, if only he is otherwise in the habit of using reason, who - when one presents him with examples of probity of purpose of steadfastness in following good maxims, of compassion and of general benevolence (involving in addition great sacrifices of advantages and comfort) - does not wish that he too might be so disposed” (4:454). The gauge-theoretic framework makes precise what Kant is describing here. The scoundrel’s recognition is not a sentimental wish but a structural inevitability: any exercise of reason that is sufficiently developed to assess practical claims at all is an exercise of reason that constitutively aims toward objectivity, and objectivity, as the present framework has argued, just is invariance under locally varying evaluative re-framings. The scoundrel’s reason, in gauge-theoretic terms, inherently couples to the connection: it cannot evaluate at all without engaging the structure that enforces cross-frame coherence. The examples of “probity of purpose” and “steadfastness in following good maxims” that Kant describes are cases in which the connection’s field strength is visibly high, cases where the agent’s orientation toward the moral law generates dynamical moral content that is palpable even to an observer who does not share the agent’s local evaluative frame. The scoundrel recognizes this field strength not because he has been morally persuaded but because his reason, insofar as it functions as reason, is sensitive to gauge-invariant content. He may act against the connection, i.e., he may privilege his local frame and refuse the demands of cross-frame coherence, but his reason cannot fail to register the objective standing (and field strength) of the structure he is defying. The “wish” Kant describes is the phenomenology of reason’s constitutive coupling to the good will field.

Why the good will “shines” independently of success. The curvature or field strength of the connection is a property of the connection’s internal structure (the pattern of cross-agent comparisons and structure of rational pressure across cases it encodes) not of the outcomes those comparisons produce in any particular case. An agent may act from the good will (may orient their practical reason toward the moral law, generating locally frame-dependent maxims whose invariant content is nonzero) and yet fail to achieve the desired outcome due to bad luck, opposition, or incapacity. The field strength is unaffected: the objective ethical content the agent’s will carries is not diminished by the failure of the outcome (by ethical matter fields failing to couple to this connection), because it was never constituted by the outcome in the first place. Indeed, the field strength governs which transitions are *rationally mandated*. The fact that the world fails to cooperate (that the mandated transitions don’t occur) doesn’t alter the structure of rational pressure any more than a charged particle failing to move (because it’s held in place by a constraint) alters the electromagnetic field it’s sitting in. The good will’s “shining” is intrinsic to its geometry/dynamical content as a connection, not derived from the coupling of this field to particular ethical matter fields (i.e., from its effects on particular ethical considerations).

The identification developed in this section completes the paper’s central argument: the ethical local invariance requirement, derived from OI’s own logic

(and supported by the symmetry bootstrapping argument), necessitates an ethical connection field structure. The ethical local invariance requirement can be identified with the moral law, and the good will with the connection field this requirement necessitates. The good will's unique normative status (including its unconditional value) is due to its role as a connection field and its gauge-invariant character. The remaining task is to examine where the analogy between gauge theory and moral philosophy breaks down and to show that these breakdowns are themselves philosophically productive (§4).

4. Disanalogies as Opportunities

The analogy between gauge theory and moral philosophy developed in §§2–3 is, I have argued, more than a metaphor: the structural logic of local invariance, connection, and curvature applies to both domains with genuine formal precision. But the analogy is not identity. Physics and ethics differ in ways that, if left unexamined, could be mistaken for defects in the framework. The purpose of this section is to show that the principal disanalogies are not defects but *productive*: each identifies a feature of the ethical domain that the gauge-theoretic framework makes newly visible and newly tractable. A framework that merely reproduced gauge theory in ethical clothing would be suspicious, a sign that the analogy was being forced rather than discovered. What we want, and what we find, is a framework in which the structural parallels illuminate and the structural differences *generate philosophical work*. Four disanalogies merit sustained attention.

4.1 The non-uniqueness of the connection

In physics, the gauge group of a theory is determined by a combination of theoretical considerations and empirical evidence. Electromagnetism has gauge group $U(1)$; the electroweak theory has $SU(2) \times U(1)$; quantum chromodynamics has $SU(3)$. For a given gauge group, the connection (gauge field) is the unique kind of structure that enforces local invariance, though its specific *configuration* (the particular field values at each point) varies across solutions.

In the ethical case, the situation is different in a philosophically consequential way. The argument of §2 shows that local evaluative invariance *necessitates* a connection, i.e., some structure that translates evaluative content across locally varying frames. And the argument of §3 shows that the Kantian good will fits the structural profile of the connection. But the argument does not show that the good will is the *only* structure that could fill this role.

Consider the candidates. The Kantian good will, the orientation of practical reason toward acting on universalizable maxims, is one way of enforcing cross-frame coherence. But a Confucian account centered on *rén* (benevolence/humaneness), understood as the cultivated sensitivity to the relational demands of one's position, could also plausibly serve as a connection: it specifies how an agent embedded in a particular web of relationships translates the

ethical demands of that position into content recognizable from other relational positions. An Aristotelian account centered on *phronesis* (practical wisdom), understood as the capacity to perceive the ethically salient features of a situation and respond appropriately, could serve a similar role: it is the structure that enables an agent with a particular character and history to generate ethical content that is not prisoner to that particular standpoint.

Is this non-uniqueness a problem? I argue it is not, and that it is, in fact, one of the framework’s most illuminating results. Here is why.

(i) Non-uniqueness as a diagnostic tool. The framework converts a familiar but often murky debate, “Is Kantian ethics or Confucian ethics or virtue ethics correct?” into a precise structural question: *do these traditions offer the same connection (expressed in different local frames) or genuinely different connections?* If the good will, *rén*, and *phronesis* are different local expressions of the same underlying connection (related by gauge transformations, as I suggested in §3.2) then the traditions’ surface-level disagreements are representational, not substantive, and the framework explains *why* they converge in practice across a wide range of cases while diverging in their conceptual vocabulary. If, on the other hand, they are genuinely different connections, then they will generate different curvatures, i.e., different gauge-invariant ethical content, and the framework predicts *where* and *how* the traditions will diverge in their substantive ethical verdicts. In either case, the framework transforms a traditionally intractable metaethical question into one with a determinate formal structure and, in principle, testable implications.

(ii) Non-uniqueness as a constraint on moral realism. If multiple genuinely distinct connections are compatible with the local invariance requirement, this is a substantive result about the structure of the moral domain: the requirement of objective evaluation constrains the *form* a connection must take (it must satisfy the transformation law of equation: $C_{ij} \mapsto g_x(j) \cdot C_{ij} \cdot g_x(i)^{-1}$) without fixing its content uniquely. This is a precise formulation of a form of *moral underdetermination*: the structural requirements of objectivity are necessary but not sufficient to single out a unique moral framework. This result bears on debates about moral realism: it suggests that if moral realism is true, its truth involves more than the structural requirements that local invariance imposes. There must be additional considerations (analogous to the empirical evidence that selects among gauge groups in physics) that determine which connection is actual. What those additional considerations might be, e.g., experience, reflective equilibrium, the internal coherence of a tradition, something else entirely, is a question the framework opens rather than answers, but it opens it with a precision that the pre-formal debate lacked.

(iii) Non-uniqueness and the space of ethical theories. In physics, the classification of possible gauge theories (the “landscape” of quantum field theories with different gauge groups) is itself a subject of deep theoretical investigation. The ethical analog would be a systematic classification of possible connections compatible with local evaluative invariance: a map of the “space of

possible ethical theories” characterized by their gauge-theoretic structure. Different connections would correspond to different ways of organizing the ethical landscape, each satisfying the structural requirements of objectivity but differing in the invariant content they generate. Charting this space is a natural program for future work; the framework provides the mathematical vocabulary in which to conduct it.

4.2 Self-referentiality: the connection as internal to practical reason

In physics, the gauge field is a structure that exists “alongside” the matter fields it couples to. The electromagnetic field is not constituted by the charged particles whose interactions it mediates; it is an independent dynamical entity with its own degrees of freedom. The gauge field and the matter fields are ontologically distinct, even if they are dynamically coupled.

In the ethical case, the situation is strikingly different. The connection (the good will, or whatever structure fills its role) is not an entity that exists alongside the agents whose frames it relates. It is an *orientation of the agents’ own practical reason*. The agents whose evaluative frames the connection translates are the very agents who exercise, cultivate, and sustain the connection through their moral activity. The connection is, in this sense, *internal* to practical reason rather than external to it.

This self-referentiality (the fact that the ethical gauge field is constituted by the very agents it connects) is a genuine disanalogy with physics. But it is a productive one for at least two reasons.

First, it illuminates a distinctive feature of the ethical domain that the physical analogy might otherwise obscure: ethical structure is not imposed on agents from outside but is generated by the agents’ own rational activity. The good will is not a field that exists independently of rational agents and then happens to couple to them; it is what rational agents *do* when they orient their practical reason toward objectivity. This makes the ethical case, in one respect, more intimate than the physical case: the connection is not discovered but *enacted*. The moral community does not find a pre-existing gauge field; it constitutes one through the activity of practical reason.²³

Second, the self-referentiality has consequences for the dynamics of the moral community that have no direct physical analog. In physics, the gauge field responds to matter through coupling but is not itself a mode of the matter’s activity. In ethics, the “strength” of the connection, i.e., how effectively it

²³This self-constitutive character of the ethical connection has an interesting resonance with Korsgaard’s constitutivist project (1996, 2009), which argues that moral normativity is grounded in the self-constituting activity of rational agents. The gauge-theoretic framework offers a different route to a structurally related conclusion: the connection is “self-constituted” not in the sense that individual agents create it through acts of autonomous self-legislation, but in the sense that it is constituted by the pattern of cross-agent rational relations that the agents themselves sustain.

enforces cross-frame coherence, depends on the agents' own moral engagement: their willingness to reason from principle, their cultivation of the capacity to recognize obligations across evaluative differences, their sustained attention to the demands of impartiality. A moral community in which agents fail to exercise practical reason toward objectivity is a community in which the connection weakens, not because an external field has decayed, but because the agents have ceased to do the work that constitutes it. This gives formal expression to a thought that is central to the Kantian tradition: moral life requires ongoing effort; the good will is not a possession but an *activity*.²⁴

4.3 The absence of a variational principle

In physics, the dynamics of the gauge field are selected by a variational principle: among all possible configurations of the field, the physical configuration is the one that extremizes (typically minimizes) an action functional. The Yang-Mills action, for instance, is a functional of the curvature (roughly, the integral of the curvature's "squared magnitude" over spacetime) and the field equations that govern the gauge field's behavior are obtained by demanding that this action be stationary under small variations. The action principle thus provides a *selection criterion*: it picks out the physically realized configuration from the space of all possible configurations.

The ethical framework, as developed here, has no obvious analog of the action principle. The local invariance requirement tells us that a connection must exist and constrains its transformation properties; the curvature characterizes its gauge-invariant content; but nothing in the argument selects which particular connection, among those satisfying the structural requirements, is the "correct" one.

This disanalogy is real, but I contend it points toward a research opportunity rather than a gap.

What would an ethical action principle look like? An action principle, if one existed, would be a functional of the curvature \mathcal{E} that, when extremized, selects the "dynamically favored" connection. In the ethical case, this would mean: a criterion for evaluating connections (moral frameworks) by examining the gauge-invariant content (curvature) they generate, and selecting the connection that optimizes this criterion.

Several candidates suggest themselves, each with recognizable philosophical pedigree. One could consider a *flatness principle*: the "best" connection is the one closest to flatness, i.e., the one whose curvature is minimized. This would correspond to a moral framework in which cross-agent evaluative relations are as globally coherent as possible, with minimal irreducible discrepancy. This has an

²⁴Kant himself insists that the good will is a matter of *Gesinnung* (fundamental orientation or disposition) rather than of isolated acts: it is the sustained orientation of practical reason toward the moral law. The gauge-theoretic framework formalizes this by locating the connection's existence in the ongoing pattern of cross-agent translation and objective deliberative motion, not in any single evaluative act.

affinity with certain forms of moral universalism that treat evaluative diversity as, ideally, resolvable.

Alternatively, one could consider a principle that *maximizes* certain curvature invariants, corresponding to a moral framework that prizes the richness and irreducibility of cross-agent moral structure, a framework in which the moral community's evaluative diversity is not a defect to be eliminated but a source of genuine moral content. This would have an affinity with forms of moral pluralism.

A third possibility (perhaps the most interesting) would be a principle analogous to the Yang-Mills action, which does not minimize or maximize curvature but selects configurations in which the curvature is *self-consistent*: in which the connection and the curvature it generates are in a kind of reflective equilibrium. This would correspond to a moral framework selected not for flatness or richness per se but for internal coherence: the moral structure the framework generates is the moral structure it requires.²⁵

I do not advocate for any particular candidate here. The point is that the absence of an action principle is not a lacuna in the framework but an *open parameter*, one that the framework identifies precisely and that different ethical traditions might fill differently. The framework provides the formal vocabulary in which competing action principles can be stated, compared, and evaluated.

4.4 The normative character of the connection

A final disanalogy concerns the mode of existence of the connection itself. The gauge field in physics *describes* what is the case: it is a component of the physical world whose configuration is a matter of empirical fact. The ethical connection, by contrast, *prescribes*: it specifies not what agents' evaluative relations *are* but what they *ought to be* if objective evaluation is to be possible.

This is a genuine difference, and it would be a mistake to minimize it. But it is also a difference that the framework accommodates naturally rather than one that threatens it.

The key observation is that the derivation of the connection in §2 is *conditional*: *if* objective ethical evaluation is to be locally gauge-invariant, *then* a connection with specific structural properties is required. The argument does not assert that the connection exists as a brute metaphysical fact; it shows that it is a *condition of the possibility* of objective ethical evaluation. This conditional structure is shared with Kant's own arguments in his critical philosophy, which establish not that their conclusions are true simpliciter but that they are necessary conditions of the relevant domain's possibility.²⁶

²⁵The resonance with Rawlsian reflective equilibrium is deliberate. The gauge-theoretic framework suggests that reflective equilibrium, understood as a coherence condition on the relationship between moral principles and their consequences, might have a natural formalization as a condition on the relationship between a connection and its curvature; see Rawls (1971). Developing this suggestion is beyond the scope of the present paper but represents a natural direction for future work.

²⁶The conditional character of the argument is a strength, not a weakness. It implies the framework is compatible with a range of metaethical positions. A moral realist can read the

The normative character of the connection does, however, open a distinctive question that has no physical analog: the question of *compliance*. In physics, the gauge field is what it is; matter couples to it necessarily, not voluntarily. In ethics, agents can fail to orient their practical reason toward objectivity; they can act from inclination rather than duty, privilege their local frame, refuse the demands of cross-frame coherence. The connection is, as noted in §4.2, constituted by agents' rational activity, and that activity can falter.

The framework gives this familiar moral-philosophical observation a new formal articulation. An agent who fails to act from the good will is, in gauge-theoretic terms, an agent who disrupts the connection: who acts as though no translation between local frames were needed, as though their own evaluative frame were privileged. The result is not merely a “wrong action” in the sense of a bad outcome; it is a *structural defect* in the moral community's gauge field, a defect whose consequences, like the consequences of a discontinuity in a physical gauge field, propagate beyond the individual agent and affect the coherence of the entire evaluative network.

This suggests that the moral significance of individual moral failure—the sense in which one agent's departure from duty affects the moral community as a whole, not merely the agent themselves—has a natural gauge-theoretic characterization. It is not merely that bad actions produce bad consequences (a frame-dependent assessment) but that failures of the good will compromise the connection itself, degrading the community's capacity for objective ethical evaluation and to realize an objective ethical order. The “damage” is to the gauge structure, not only in the outcomes.

Summary

The four disanalogies examined in this section (non-uniqueness, self-referentiality, the absence of a variational principle, and the normative character of the connection) are not defects in the gauge-theoretic framework applied to ethics. They are structural features of the ethical domain that the framework makes formally visible and philosophically tractable. Each opens a research direction: the classification of possible ethical connections, the dynamics of moral communities as self-constituting gauge systems, the search for ethical action principles, and the formal characterization of moral failure as gauge-structural defect. That the analogy with physics is *imperfect* is precisely what makes it *productive*: the points of disanalogy mark the places where the ethical domain has structure that physics does not, and the gauge-theoretic vocabulary provides new tools for exploring and articulating that structure.

connection as a structure that objectively exists and whose existence is established (conditionally) by the argument. A constructivist can read it as a structure that rational agents constitute through the activity of practical reason. An anti-realist can read the argument as showing what the structure of morality *would have to be* if moral objectivity were achievable, thereby illuminating what is at stake in denying it.

5. Conclusion

This paper has argued that the transition from ethical global to local invariance in the objectivity-as-invariance framework is not merely a formal refinement but a philosophically consequential step: it forces the existence of new structure (an ethical connection field), identifies its gauge-invariant content (its curvature/field strength), and yields a precise characterization of the unconditional value of the good will as the evaluative content that survives every local re-framing.

The argument's arc can be summarized in four claims. First, if objective ethical evaluation must be invariant not only under global permutations of agents but under locally varying transformations of agents' evaluative frames, then ethical evaluation requires a connection field: a structure that specifies how to translate ethical content across differently calibrated standpoints (§2). Second, this connection field is not itself fully gauge-invariant; its objective contribution to the ethical landscape resides in its curvature or dynamical content: the irreducible, frame-independent ethical structure that the connection carries and that makes frame-independent evaluation possible (§2.4). Third, the Kantian good will fits the structural profile of this connection field: it enforces coherence across locally varying frames, is necessitated by the demand for practical objectivity, varies in its local expression while preserving its structural function, and generates genuine dynamical ethical content (the structure of rational pressure across cases) through its gauge-invariant curvature/field strength (§3). Fourth, the disanalogies between the physical and ethical cases (the non-uniqueness of the connection, the self-referentiality of the ethical gauge field, the absence of a variational principle, and the normative character of the connection) are not defects but productive features that identify distinctive structures of the ethical domain and open determinate directions for future work (§4).

Two programmatic remarks are in order.

The Kant-Einstein parallel. Kant famously remarked that Rousseau was the Newton of the moral world: the thinker who first discovered, beneath the apparent diversity of human moral life, a deep underlying order governed by universal law.²⁷ That is, Kant seemed to think that Rousseau found the simple universal law of autonomy that explains the structure of the moral order, just as Newton found simple universal laws of motion that explain the structure of the physical order. If the framework developed here is on the right track, then there is a sense in which Kant can be said to be the Einstein of the moral world. In particular, there is a deep methodological parallel between Kant and Einstein:

Before Einstein, physics had symmetries, but they were treated as discovered properties of particular theories. Newtonian mechanics happened to be Galilean invariant; Maxwell's equations happened to be Lorentz invariant. The symmetries were consequences, not premises. Einstein inverted this relationship: he

²⁷The remark appears in Kant's handwritten notes (the *Bemerkungen*, Ak. 20:58–59). Kant writes that just as Newton found simplicity and order where previously there had been complexity and discord, Rousseau discovered the deeply hidden nature of the human being and the concealed law that justifies providence.

took the symmetry requirement (the principle of relativity, the invariance of the speed of light) as foundational — as an a priori constraint on admissible physical theories — and derived the dynamics (special relativity, and then general relativity) as consequences. The symmetry principle came first; the physics followed. This was not merely a different way of arriving at the same theory; it was a fundamentally different conception of the relationship between symmetry and dynamics. It transformed physics from a discipline that discovers symmetries in its theories to a discipline that uses symmetries to constrain which theories are admissible.

Kant did exactly this for ethics. Before Kant, moral philosophy had impartiality requirements (e.g., the Golden Rule, natural law precepts, and utilitarian maximization), but these principles were defended on their own terms, and universality was a feature they exhibited. Kant inverted this: he took the universalizability requirement (the categorical imperative) as *foundational* — as an a priori constraint on admissible maxims — and derived the moral content as what survives the constraint. The symmetry principle came first; the ethics followed. Admissible maxims are those that are invariant under universalization, just as admissible physical theories are those that are invariant under the relevant coordinate transformations.

The depth of the parallel extends to the specific kind of symmetry principle each thinker invoked. The move from Newton to Einstein is a move from global symmetries and conservation laws to local symmetries and gauge structure (general relativity's different coordinate transformations at different points, compensated by the dynamical metric as a connection). This move has a structural counterpart in the move from more standard moral requirements of impartiality (global symmetries of practical reason) to Kant's own universalizability requirement (a local symmetry requirement that necessitates a connection). The good will, understood as a gauge connection, is not merely the moral analog of a Newtonian conserved quantity; it is the moral analog of a gauge field (such as general relativity's dynamical metric field), a richer, more dynamical structure whose objective content (field strength) emerges from the requirement that ethical evaluation be coherent across locally varying evaluative frames, just as the gauge field's dynamical content emerges from the requirement that physics be coherent across local symmetry transformations. That the foundations of Kant's moral philosophy (in particular, the characteristic relationship between the moral law and the good will) can be faithfully formalized using gauge theory is, at a minimum, a striking convergence.²⁸

²⁸The suggestion that Kant was informally articulating a structure that mathematics has since formalized as local gauge symmetry may seem anachronistic. But there is a well-established precedent for this kind of retroactive formalization in the history of philosophy. Game-theoretic structures (e.g., the prisoner's dilemma, the stag hunt, coordination games) have been read back into Hobbes's *Leviathan* and Hume's *Treatise* by scholars such as Jean Hampton (1986) and Brian Skyrms (2004), and these readings are widely regarded as illuminating rather than anachronistic precisely because the strategic structures were already operative in the philosophical arguments, awaiting a formal vocabulary adequate to make them explicit. The claim here is analogous: the structural relationship between the moral law (as a requirement of frame-independent evaluation) and the good will (as the structure

The empty formalism objection. A longstanding criticism of Kantian ethics holds that the categorical imperative is “empty,” i.e., that the universalizability test, by itself, cannot generate substantive moral content and must smuggle in material premises to yield determinate duties.²⁹ The gauge-theoretic framework suggests a novel line of response. If the good will is a connection, then its “content” is not reducible to the pointwise verdicts it generates (which are, in any case, frame-dependent); it resides in the curvature/field strength, i.e., in the gauge-invariant structure of cross-agent ethical relations and rational pressure across cases. The empty formalism objection, on this reading, mistakes the frame-dependence of the connection’s local expression for a lack of content, when in fact the content is present but lives at the level of curvature rather than at the level of individual connection components. Developing this response in full, showing precisely how the gauge-theoretic framework dissolves the objection by relocating moral content from the connection to the curvature, is a task for future work. But the framework’s resources for addressing the objection are, I believe, considerable, and the direction of the response is already visible in the formal apparatus of §2.

If the argument of this paper succeeds, its contribution is twofold. At the level of formal ethics, it shows that the logic of local gauge invariance (one of the deepest structural principles in modern physics) applies to the ethical domain with genuine precision, forcing the existence of a connection structure and characterizing its gauge-invariant content in a way that illuminates the Kantian good will. At the level of methodology, it demonstrates that the analogy between physics and ethics, when pursued with formal care and honest attention to disanalogies, is not a loose metaphor but a source of determinate philosophical results, results that could not have been obtained without the gauge-theoretic vocabulary and that open a structured research program extending from the classification of possible ethical connections to the search for ethical action principles to the formal characterization of moral failure as gauge-structural defect.

The good will, Kant tells us, shines by its own light. The gauge-theoretic framework lets us say, with a precision not available to Kant, what that light is: it is the curvature or field strength of the ethical connection, the irreducible, frame-independent structure that practical reason, oriented toward objectivity, carries into every evaluative situation. It does not borrow its significance from the local frame in which it is expressed, nor from the consequences it produces. It shines because it is the structure that makes objective ethical evaluation possible across the full diversity of human evaluative life and that encodes the objective action-guiding content of ethics. It shines unconditionally because its light is gauge-invariant.

that enforces coherence across locally varying frames) is already present in the *Groundwork’s* architecture. What gauge theory provides is not a foreign imposition but a formal language that makes visible a structure Kant was already working with, much as game theory made visible the strategic structures Hobbes and Hume were already working with.

²⁹The objection traces to Hegel’s critique in the *Philosophy of Right* (§135) and has been developed in various forms by Mill and Schopenhauer among others.

References

- [1] Hampton, J. (1986). *Hobbes and the social contract tradition*. Cambridge University Press.
- [2] Hegel, G. W. F. (1991). *Elements of the philosophy of right* (A. W. Wood, Ed.; H. B. Nisbet, Trans.). Cambridge University Press. (Original work published 1821).
- [3] Kant, I. (1785/1998). *Groundwork of the metaphysics of morals* (M. Gregor, Trans.). Cambridge University Press. Cited by Akademie edition pagination.
- [4] Kant, I. *Bemerkungen* (Handwritten notes). Akademie edition, vol. 20.
- [5] Kogut, J. B. (1979). An introduction to lattice gauge theory and spin systems. *Reviews of Modern Physics*, 51(4), 659–713.
- [6] Korsgaard, C. M. (1996). *The sources of normativity*. Cambridge University Press.
- [7] Korsgaard, C. M. (2009). *Self-constitution: Agency, identity, and integrity*. Oxford University Press.
- [8] Noether, E. (1918/2011). Invariant variational problems. In Y. Kosmann-Schwarzbach & B. E. Schwarzbach (Eds.), *The Noether theorems* (pp. 3–22). Springer.
- [9] Rawls, J. (1971). *A theory of justice*. Harvard University Press.
- [10] Sanchez Borboa, S. D. J. (manuscript A). Symmetry and the form of ethical objectivity.
- [11] Sanchez Borboa, S. D. J. (manuscript B). Interacting ethical fields: The coupled dynamics of the good, the right, and moral worth.
- [12] Skyrms, B. (2004). *The stag hunt and the evolution of social structure*. Cambridge University Press.
- [13] Smit, H. (2009). Kant on apriority and the spontaneity of cognition. *In Metaphysics and the good: Themes from the philosophy of Robert Merrihew Adams*, pp. 188–251.
- [14] Smit, H. (forthcoming). Kant’s theory of cognition
- [15] Wilson, K. G. (1974). Confinement of quarks. *Physical Review D*, 10(8), 2445–2459.