

COHOMOLOGICAL CALIBRATION AND CURVATURE CONSTRAINTS ON PRODUCT MANIFOLDS: A TOPOLOGICAL LOWER BOUND

ALEXANDER PIGAZZINI, MAGDALENA TODA

ABSTRACT. We establish a quantitative relationship between mixed de Rham classes and the geometric complexity of metric connections with totally skew torsion on product manifolds where both factors are compact oriented surfaces. For any cohomologically calibrated connection ∇^C whose torsion T has pure bidegree with respect to the product decomposition and whose harmonic projection represents a non-trivial mixed class $[\omega]$, we prove that on a non-empty open subset $\mathcal{V} \subset M$,

$$\dim(\mathfrak{ho}_p^{\text{off}}(\nabla^C)) \geq r^\sharp := \text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}}) - \dim \mathcal{K},$$

with \mathcal{K} an intrinsically defined obstruction space. The bound is a topological invariant under metric deformations preserving the parallel-form strata and provides an obstruction to reducible holonomy. Counterexamples show the hypothesis is optimal.

Keywords: Connections with Torsion, Product Manifolds, Holonomy Algebra, De Rham Cohomology, Hodge Theory, Künneth Decomposition, Curvature Tensor

MSC Classification 2020: 53C05, 53C07, 53C29, 58A14, 57R19, 15A69

1. INTRODUCTION

Let $M = M_1 \times M_2$ be a compact oriented product manifold endowed with a product metric $g = g_1 \oplus g_2$, inducing at each point an orthogonal splitting $T_p M = V_1 \oplus V_2$. We introduce a topological invariant, the *mixed tensor rank* of a cohomology class $[\omega] \in H^3(M; \mathbb{R})$, and show that—after subtracting a correction term that captures the part of the class absorbed by parallel forms—it provides a lower bound for the dimension of the off-diagonal holonomy subspace generated by any *cohomologically calibrated* metric connection ∇^C with totally skew torsion of pure bidegree.

A subtle but crucial point is that the subspaces spanned by a minimal harmonic decomposition are uniquely determined only under topological restrictions on the factors. Consequently, our main result holds for *admissible* product manifolds (Definition 2.2), i.e., where both factors are compact oriented surfaces.

For any admissible M , any cohomologically calibrated connection ∇^C with torsion of pure bidegree satisfies, on a non-empty open subset $\mathcal{V} \subset M$,

$$\dim(\mathfrak{ho}_p^{\text{off}}(\nabla^C)) \geq r^\sharp := \text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}}) - \dim \mathcal{K},$$

where the obstruction space \mathcal{K} is intrinsically defined via intersections of the mixed factor spaces $\mathcal{V}_{2,1}, \mathcal{V}_{1,2}$ with the parallel-form strata $\mathcal{P}_k(M_i)$. The bound is invariant under metric deformations preserving these strata.

When $r^\sharp > 0$ the holonomy representation is forced to be irreducible across the product splitting. The examples in Section 6 illustrate both a non-trivial bound ($\Sigma_g \times T^2$ with a pure bidegree class) and a trivial bound ($T^2 \times T^2$, where the non-pure bidegree and full parallelism of the torsion yield $r^\sharp = 0$). Remark 2.4 shows by explicit counterexample that without the admissibility hypothesis the bound fails to be well-defined.

The pure bidegree condition on the torsion—namely, that T lies entirely in $\Gamma(\Lambda^2 V_1^* \otimes V_2^*)$ or entirely in $\Gamma(V_1^* \otimes \Lambda^2 V_2^*)$ —admits a transparent geometric interpretation in the product setting: since $\Lambda^2 V_i^*$ is one-dimensional for a surface, a 3-form of pure bidegree $(2, 1)$ necessarily has the form $\text{vol}_{M_1} \wedge \beta$ for a 1-form β on M_2 , so the torsion couples the two factors through the area form of exactly one of them. This condition arises naturally in Kaluza–Klein constructions on products of surfaces, where the torsion 3-form is built as the wedge product of the curvature 2-form of one factor with a connection 1-form on the other: since the curvature of a surface is necessarily proportional to its area form, the resulting torsion has pure bidegree by construction. We nevertheless conjecture (Conjecture 7.2) that the theorem extends to torsion forms of mixed bidegree, i.e. when both $T^{2,1}$ and $T^{1,2}$ are simultaneously non-vanishing.

2. COHOMOLOGICALLY CALIBRATED METRIC CONNECTIONS WITH SKEW TORSION

Throughout, ∇^{LC} denotes the Levi-Civita connection of g . We consider metric connections ∇^C with totally skew torsion T , i.e.

$$\nabla^C g = 0, \tag{2.1}$$

$$T(X, Y, Z) := g(T(X, Y), Z) \in \Omega^3(M), \tag{2.2}$$

and write

$$\nabla^C = \nabla^{LC} + K, \tag{2.3}$$

$$K_{XYZ} = \frac{1}{2} T_{XYZ}. \tag{2.4}$$

Thus $T^\flat := T$ is a 3-form. The curvature R^C of ∇^C can be expressed in terms of R^{LC} , $\nabla^{LC} T$, and quadratic torsion terms; in abstract index notation,

$$R_{XY}^C = R_{XY}^{LC} + \frac{1}{2} ((\nabla_X^{LC} T)(Y, Z, \cdot)^\sharp - (\nabla_Y^{LC} T)(X, Z, \cdot)^\sharp) + \frac{1}{4} Q_T(X, Y)Z, \tag{2.5}$$

where Q_T is bilinear and quadratic in T (see, e.g., [4] or standard torsion-curvature formulas). We shall use (2.5) only qualitatively: mixed components of T produce mixed components of R^C through both the $\nabla^{LC} T$ and $T * T$ terms.

The product structure $TM = V_1 \oplus V_2$ induces a bigrading on $\Omega^3(M)$: every 3-form decomposes pointwise as $T = T^{2,1} + T^{1,2}$, where $T^{2,1} \in \Gamma(\Lambda^2 V_1^* \otimes V_2^*)$

and $T^{1,2} \in \Gamma(V_1^* \otimes \Lambda^2 V_2^*)$ (components of type $(3,0)$ and $(0,3)$ vanish since $\dim V_i = 2$). We say that T has *pure bidegree* if one of $T^{2,1}, T^{1,2}$ vanishes identically.

Definition 2.1 (Metric cohomological calibration). A metric connection ∇^C with totally skew torsion T , where the torsion 3-form represents a non-trivial mixed cohomology class $[\omega]$ via its harmonic projection ($T^b = \omega_h + d\alpha + \delta\beta$, i.e. $[\omega] := [\omega_h]$), is called *cohomologically calibrated*. The class is *mixed* if it lies outside the natural image of $H^3(M_1) \oplus H^3(M_2)$ under the Künneth isomorphism:

$$H^3(M) \cong \bigoplus_{p+q=3} H^p(M_1) \otimes H^q(M_2). \quad (2.6)$$

Since M is compact and oriented, Hodge theory applies to g : each cohomology class has a unique harmonic representative. We denote by ω_h the harmonic representative of $[\omega]$ with respect to g .

For a thorough treatment of Hodge theory and cohomology on complex manifolds, see [6].

2.1. Intrinsic Definition of Mixed Cohomology Subspaces. We now address a subtle but fundamental point: the subspaces spanned by the factors in a minimal harmonic decomposition are not uniquely determined in general. To obtain a well-defined geometric invariant we must impose conditions on the topology of the factors.

Definition 2.2 (Admissible Product Manifolds). A compact oriented product manifold $M = M_1 \times M_2$ is called *admissible* if it satisfies the following:

both factors M_1, M_2 are compact oriented surfaces.

Lemma 2.3 (Uniqueness of Mixed Subspaces). *Let M be an admissible product manifold and $[\omega] \in H^3(M; \mathbb{R})$ a non-trivial mixed class. Then the subspaces*

$$\mathcal{V}_{2,1} \subset \mathcal{H}^2(M_1) \otimes \mathcal{H}^1(M_2), \quad (2.7)$$

$$\mathcal{V}_{1,2} \subset \mathcal{H}^1(M_1) \otimes \mathcal{H}^2(M_2) \quad (2.8)$$

spanned by any minimal harmonic decomposition of the mixed Künneth components of $[\omega]$ are intrinsically defined, i.e. they depend only on $[\omega]$ and not on the particular minimal decomposition chosen.

Proof. For a compact oriented surface F^2 we have $b_2(F) = 1$ and $b_1(F) = 2g(F)$. Let vol_F be the unique (up to scale) g_F -harmonic 2-form. Then

$$\mathcal{H}^2(F) = \mathbb{R} \cdot \text{vol}_F.$$

Consequently any non-zero $T \in \mathcal{H}^2(M_1) \otimes \mathcal{H}^1(M_2)$ has rank one and can be written uniquely as

$$T = \text{vol}_{M_1} \otimes \beta, \quad \beta \in \mathcal{H}^1(M_2).$$

Indeed, if $\text{vol}_{M_1} \otimes \beta = \text{vol}_{M_1} \otimes \beta'$, then $\text{vol}_{M_1} \otimes (\beta - \beta') = 0$ and, because vol_{M_1} is non-zero as an element of $\mathcal{H}^2(M_1)$, we must have $\beta = \beta'$. Thus the one-dimensional subspace $\text{span}\{\beta\}$ is uniquely determined by T , and

$$\mathcal{V}_{2,1} = \mathcal{H}^2(M_1) \otimes \text{span}\{\beta\}$$

is intrinsically defined. The same argument applied symmetrically yields the intrinsic definition of $\mathcal{V}_{1,2}$. \square

Remark 2.4 (Failure of invariance in general). Invariance fails without the admissibility hypothesis. For instance, on $M = T^3 \times T^3$ we have $b_2(T^3) = 3$ and $b_1(T^3) = 3$. The mixed class

$$[\omega] = (dx \wedge dy) \otimes dz + (dx \wedge dz) \otimes dy \in H^2(T^3) \otimes H^1(T^3)$$

has rank 2. It admits two distinct minimal decompositions

$$\begin{aligned} \omega_h^{2,1} &= (dx \wedge dy) \otimes dz + (dx \wedge dz) \otimes dy \\ &= \frac{1}{2} [(dx \wedge dy + dx \wedge dz) \otimes (dz + dy) \\ &\quad + (dx \wedge dy - dx \wedge dz) \otimes (dz - dy)], \end{aligned}$$

which generate different subspaces $\mathcal{V}_{2,1}$ and $\mathcal{V}'_{2,1}$. Consequently the integer $\dim(\mathcal{V}_{2,1} \cap (\mathcal{P}_2(M_1) \otimes \mathcal{P}_1(M_2)))$ depends on the choice, contradicting the metric-invariance.

2.2. Mixed obstruction spaces. Having established the intrinsic nature of the mixed factor spaces in Lemma 2.3 for admissible product manifolds (Definition 2.2), we now define the obstruction kernels.

Definition 2.5 (Mixed obstruction spaces). For an admissible product manifold M and a mixed class $[\omega]$ we define

$$\begin{aligned} \mathcal{V}_{2,1} &:= \text{span}\{\alpha_i \otimes \beta_j\} \subset \mathcal{H}^2(M_1) \otimes \mathcal{H}^1(M_2), \\ \mathcal{V}_{1,2} &:= \text{span}\{\tilde{\alpha}_j \otimes \tilde{\beta}_i\} \subset \mathcal{H}^1(M_1) \otimes \mathcal{H}^2(M_2), \end{aligned}$$

where the spans are taken over any minimal harmonic decomposition of the mixed Künneth components of $[\omega]$; by Lemma 2.3 these spans are independent of the particular choice.

The *mixed obstruction spaces* are the intersections

$$\begin{aligned} \ker \Psi_p &:= \mathcal{V}_{2,1} \cap (\mathcal{P}_2(M_1) \otimes \mathcal{P}_1(M_2)), \\ \ker \tilde{\Psi}_p &:= \mathcal{V}_{1,2} \cap (\mathcal{P}_1(M_1) \otimes \mathcal{P}_2(M_2)), \end{aligned}$$

and we set $\mathcal{K} := \ker \Psi_p \oplus \ker \tilde{\Psi}_p$. The notation $\ker \Psi_p$ is anticipatory (see Remark 2.6); the intersection itself is independent of the point $p \in M$.

Remark 2.6 (Anticipatory notation). The notation “ $\ker \Psi_p$ ” is anticipatory: in the proof of Theorem 5.2, we construct obstruction maps $\Psi_p : \mathcal{V}_{2,1} \rightarrow \mathfrak{so}(T_p M)$ and $\tilde{\Psi}_p : \mathcal{V}_{1,2} \rightarrow \mathfrak{so}(T_p M)$ via covariant derivatives of the torsion form. A key step in the proof is to verify that the kernels of these maps coincide with the intersections defined above. The induced quotient

maps $V_{2,1}/\ker \Psi_p \rightarrow \mathfrak{so}(T_p M)$ are then shown to be injective on a non-empty open subset, establishing the lower bound.

3. OFF-DIAGONAL CURVATURE AND HOLONOMY COMPLEXITY

Let $P_i : T_p M \rightarrow V_i$ be the orthogonal projections of the product splitting.

Definition 3.1 (Off-diagonal curvature). For a curvature tensor R , define its off-diagonal component by

$$R_{\text{off}}(X, Y)Z := P_1 R(X, Y)(P_2 Z) + P_2 R(X, Y)(P_1 Z). \quad (3.1)$$

Definition 3.2 (Off-diagonal holonomy component). Let (M, g) be a Riemannian product manifold with orthogonal splitting $T_p M = V_1 \oplus V_2$, and let ∇^C be a metric connection on TM with totally skew-symmetric torsion.

Let $P_i : T_p M \rightarrow V_i$ be the orthogonal projections and let $\Pi_{\text{off}} : \mathfrak{so}(T_p M) \rightarrow \mathfrak{so}(T_p M)$ be the linear projection onto off-diagonal endomorphisms:

$$\Pi_{\text{off}}(A)(Z) := P_1 A(P_2 Z) + P_2 A(P_1 Z).$$

Define the space of off-diagonal skew-symmetric endomorphisms as the image of this projection:

$$\mathfrak{so}^{\text{off}}(T_p M) := \text{Im}(\Pi_{\text{off}}) = \{A \in \mathfrak{so}(T_p M) \mid A(V_1) \subset V_2, A(V_2) \subset V_1\}.$$

The *off-diagonal holonomy component* is the intersection of the holonomy algebra with this space:

$$\mathfrak{hol}_p^{\text{off}}(\nabla^C) := \mathfrak{hol}_p(\nabla^C) \cap \mathfrak{so}^{\text{off}}(T_p M).$$

Equivalently, $\mathfrak{hol}_p^{\text{off}}(\nabla^C) = \{A \in \mathfrak{hol}_p(\nabla^C) \mid \Pi_{\text{off}}(A) = A\}$ consists of those holonomy elements that are purely off-diagonal.

Remark 3.3. Although $\mathfrak{hol}_p^{\text{off}}(\nabla^C)$ is a vector subspace of $\mathfrak{hol}_p(\nabla^C)$, it is not closed under the Lie bracket, since $[\mathfrak{so}^{\text{off}}, \mathfrak{so}^{\text{off}}] \subset \mathfrak{so}(V_1) \oplus \mathfrak{so}(V_2)$. Nevertheless, $\dim \mathfrak{hol}_p^{\text{off}}(\nabla^C) > 0$ implies that the holonomy representation is irreducible with respect to the splitting $V_1 \oplus V_2$, therefore, we have global irreducibility.

By the Ambrose–Singer theorem [3], $\mathfrak{hol}_p^{\text{off}}(\nabla^C)$ depends only on the global holonomy algebra of ∇^C and is therefore independent of the choice of local generators and of the pointwise rank of the curvature tensor. In particular, at each point $p \in M$ one has, clearly, $\dim \mathfrak{hol}_p(\nabla^C) \geq \dim \mathfrak{hol}_p^{\text{off}}(\nabla^C)$. Because ∇^C is metric, its holonomy algebra sits in $\mathfrak{so}(T_p M)$. If the action were reducible with respect to $V_1 \oplus V_2$, all holonomy endomorphisms would be block-diagonal, forcing $\mathfrak{hol}_p^{\text{off}}(\nabla^C) = \{0\}$. Hence non-trivial off-diagonal holonomy implies irreducibility.

4. MIXED RANK OF A DEGREE-3 CLASS

Write the Künneth decomposition of $H^3(M)$ as:

$$\begin{aligned} & (H^3(M_1) \otimes H^0(M_2)) \oplus (H^2(M_1) \otimes H^1(M_2)) \oplus \\ & (H^1(M_1) \otimes H^2(M_2)) \oplus (H^0(M_1) \otimes H^3(M_2)). \end{aligned} \quad (4.1)$$

Mixed components occur precisely in bidegrees $(2, 1)$ and $(1, 2)$.

Definition 4.1 (Mixed tensor rank). Let $[\omega] \in H^3(M; \mathbb{R})$. Denote by $\pi_{2,1}([\omega]) \in H^2(M_1) \otimes H^1(M_2)$ and $\pi_{1,2}([\omega]) \in H^1(M_1) \otimes H^2(M_2)$ its mixed Künneth projections. Define

$$r_{2,1} = \min \left\{ r : \pi_{2,1}([\omega]) = \sum_{i=1}^r \alpha_i \otimes \beta_i, \alpha_i \in H^2(M_1), \beta_i \in H^1(M_2) \right\}, \quad (4.2)$$

and analogously $r_{1,2}$ for $\pi_{1,2}([\omega])$. The mixed rank is

$$\text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}}) := r_{2,1} + r_{1,2}. \quad (4.3)$$

For degree 3, the tensor rank in each bidegree is well-defined and finite, and can be computed from any decomposition of the harmonic representative ω_h into wedge products of harmonic forms on the factors for the product metric g . Although ω_h depends on g , the ranks $r_{2,1}, r_{1,2}$ are purely topological (indeed, the tensor rank of an element in $H^2(M_1) \otimes H^1(M_2)$ is defined as the minimal number of simple tensors in any decomposition, hence is invariant under change of metric within the product class).

5. MAIN RESULT AND PROOF

Lemma 5.1 (Existence of points of maximal rank for covariant derivatives). *Let (M_1, g_1) be a compact oriented Riemannian manifold and let $\mathcal{H}^2(M_1, g_1)$ denote the space of harmonic 2-forms. Given a subspace $U \subset \mathcal{H}^2(M_1, g_1)$ of dimension r and its parallel part $U_{\parallel} := U \cap \mathcal{P}_2(M_1, g_1)$ of dimension s , choose a basis $\alpha_1, \dots, \alpha_r$ of U such that $\alpha_1, \dots, \alpha_s$ span U_{\parallel} .*

If $r - s \leq 1$ (condition satisfied by admissible product manifolds: U is the span of a minimal harmonic decomposition, so $\dim U = r_{2,1} \leq 1$, hence $r - s \leq 1$; see Definition 4.1, Definition 2.2 and Lemma 2.3), then there exists a non-empty open set $\mathcal{U} \subset M_1$ on which the covariant derivatives

$$(\nabla^{g_1} \alpha_{s+1}, \dots, \nabla^{g_1} \alpha_r)$$

are pointwise linearly independent; i.e.

$$\mathcal{U} = \left\{ p \in M_1 \mid \text{rank}[(\nabla^{g_1} \alpha_i)(p)]_{i=s+1, \dots, r} = r - s \right\} \neq \emptyset.$$

Moreover, if g_1 is real-analytic, then \mathcal{U} is dense in M_1 .

Proof. We proceed in three steps.

1. *Reduction to a real-analytic metric.* Real-analytic Riemannian metrics are dense in the space of smooth metrics on a compact manifold; this follows from the theorem of Greene–Jacobowitz on analytic approximation of

smooth maps (see [5]). Hence we can choose a sequence of real-analytic metrics $\{g_1^{(k)}\}_{k \in \mathbb{N}}$ converging to g_1 in the C^∞ topology. The dimension of the space of parallel 2-forms is upper-semicontinuous with respect to the metric (see [4]); therefore, after passing to a subsequence, we may assume

$$\dim \mathcal{P}_2(M_1, g_1^{(k)}) \leq \dim \mathcal{P}_2(M_1, g_1) = s \quad \text{for all } k.$$

For each k let $\alpha_i^{(k)} \in \mathcal{H}^2(M_1, g_1^{(k)})$ be the unique harmonic representative of the cohomology class $[\alpha_i]$. By the stability theorem for harmonic forms under metric deformation (Warner [10], see also Kodaira–Spencer [8]), $\alpha_i^{(k)} \rightarrow \alpha_i$ in C^∞ . Consequently the covariant derivatives converge in C^1 :

$$\nabla^{g_1^{(k)}} \alpha_i^{(k)} \longrightarrow \nabla^{g_1} \alpha_i \quad (i = s + 1, \dots, r).$$

If we prove that for each k there exists a dense open set $\mathcal{U}^{(k)} \subset M_1$ where the $\nabla^{g_1^{(k)}} \alpha_i^{(k)}$ are linearly independent, then the C^1 convergence guarantees that, for large enough k , the same linear independence holds on a non-empty open set for the original metric g_1 . Indeed, linear independence is equivalent to the non-vanishing of at least one $(r - s) \times (r - s)$ minor of the matrix of components; each minor depends continuously on the metric and the forms, so a minor that is non-zero on an open set for $g_1^{(k)}$ remains non-zero on an open set for g_1 when k is sufficiently large.

We now prove the stronger density statement under the assumption that g_1 is real-analytic

2. Admissible manifolds. Assume now that g_1 is real-analytic. Then every harmonic form is real-analytic (by elliptic regularity for the Hodge–Laplacian, see [9]), and consequently the covariant derivatives $\nabla^{g_1} \alpha_i$ are real-analytic sections of the vector bundle $E := T^*M_1 \otimes \Lambda^2 T^*M_1$.

Consider the real-analytic section

$$\Phi : M_1 \longrightarrow \Lambda^{r-s} E, \quad \Phi(p) := (\nabla^{g_1} \alpha_{s+1})(p) \wedge \cdots \wedge (\nabla^{g_1} \alpha_r)(p).$$

The set $\mathcal{U} := \{p \in M_1 \mid \Phi(p) \neq 0\}$ is exactly the set where the covariant derivatives are linearly independent; because Φ is continuous, \mathcal{U} is open. We must show that $\mathcal{U} \neq \emptyset$.

Suppose, for contradiction, that $\mathcal{U} = \emptyset$; i.e. $\Phi(p) = 0$ for every $p \in M_1$. Then at each point the sections $\nabla^{g_1} \alpha_{s+1}, \dots, \nabla^{g_1} \alpha_r$ are linearly dependent. Since the condition of rank $\leq d$ is given by the vanishing of all $(d+1) \times (d+1)$ minors, which are analytic functions, the locus where the rank is constant is an analytic subset. Let $d < r - s$ be the minimal rank and let p_0 be a point where the rank equals d . There exists a connected open neighbourhood V of p_0 on which the rank is constant and equal to d .

On V we can therefore find real-analytic functions $c_1, \dots, c_{r-s} : V \rightarrow \mathbb{R}$, not all zero, such that

$$\sum_{i=1}^{r-s} c_i(p) (\nabla^{g_1} \alpha_{s+i})(p) = 0 \quad \forall p \in V. \quad (5.1)$$

Since we have to consider admissible manifolds (Definition 2.2 and Lemma 2.3) we have two cases:

Case $r = s$: There are no non-parallel forms to consider; the lemma holds vacuously.

Case $r = s + 1$: There exists a unique non-parallel form α and a single analytic function c such that (5.1) reads $c(p)(\nabla^{g_1}\alpha)_p = 0$ for all $p \in V$. If $c(p_0) \neq 0$ for some $p_0 \in V$, then by analyticity $c(p) \neq 0$ on an open subset $V' \subset V$, whence $(\nabla^{g_1}\alpha)_p = 0$ on V' . By the unique-continuation principle for analytic sections (Aronszajn [1]; for differential forms see [2]), $\nabla^{g_1}\alpha \equiv 0$ on the whole connected component containing V' , and hence on all of M_1 . This contradicts the assumption that α is non-parallel modulo $\mathcal{P}_2(M_1, g_1)$. Therefore $c \equiv 0$ on V , contradicting the choice of the c_i 's. Hence $\mathcal{U} \neq \emptyset$. Moreover, the complement $M_1 \setminus \mathcal{U} = \{p \in M_1 : (\nabla^{g_1}\alpha)(p) = 0\}$ has empty interior: if it contained an open set V' , then by unique continuation $\nabla^{g_1}\alpha \equiv 0$ on all of M_1 , contradicting the hypothesis that α is non-parallel. Therefore \mathcal{U} is dense in M_1 .

3. *Returning to the original smooth metric.* For the original smooth metric g_1 , the non-emptiness of \mathcal{U} is immediate: since $\alpha_{s+1} \notin \mathcal{P}_k(M_1, g_1)$, the section $\nabla^{g_1}\alpha_{s+1}$ is not identically zero, and $\mathcal{U} = \{p \in M_1 : (\nabla^{g_1}\alpha_{s+1})(p) \neq 0\}$ is non-empty and open. Steps 1–2 establish the stronger conclusion that \mathcal{U} is dense when g_1 is real-analytic. \square

The non-emptiness argument holds for harmonic forms of any degree k and any smooth metric. The density statement for real-analytic metrics likewise holds for any degree k , as Steps 1–2 use only analytic approximation, semicontinuity of the dimension of parallel forms, C^1 convergence, and the identity principle for real-analytic sections. Therefore, the same results hold for $H^1(M_1, g_1)$, as well as for $H^1(M_2, g_2)$ and $H^2(M_2, g_2)$

Theorem 5.2 (Topological Lower Bound for Off-Diagonal Curvature). *Let $M = M_1 \times M_2$ be an admissible compact oriented Riemannian product manifold (Definition 2.2) with product metric $g = g_1 \oplus g_2$. Let ∇^C be a cohomologically calibrated metric connection (Definition 2.1) whose torsion T has pure bidegree with respect to the product decomposition, i.e., $T \in \Gamma(\Lambda^2 V_1^* \otimes V_2^*)$ or $T \in \Gamma(V_1^* \otimes \Lambda^2 V_2^*)$ and assume further that the harmonic representative $[\omega_h]$ is a non-trivial mixed class in the Künneth decomposition of $H^3(M)$.*

Define the mixed factor subspaces

$$\mathcal{V}_{2,1} := \text{span}\{\alpha_i \otimes \beta_i\} \subset \mathcal{H}^2(M_1) \otimes \mathcal{H}^1(M_2),$$

$$\mathcal{V}_{1,2} := \text{span}\{\tilde{\alpha}_j \otimes \tilde{\beta}_j\} \subset \mathcal{H}^1(M_1) \otimes \mathcal{H}^2(M_2),$$

where the spans are taken over any minimal harmonic decomposition of the mixed Künneth components of $[\omega]$; in the admissible setting these spans are independent of the particular choice (Lemma 2.3).

Define the obstruction kernels

$$\ker \Psi_p := \mathcal{V}_{2,1} \cap (\mathcal{P}_2(M_1) \otimes \mathcal{P}_1(M_2)), \quad \ker \tilde{\Psi}_p := \mathcal{V}_{1,2} \cap (\mathcal{P}_1(M_1) \otimes \mathcal{P}_2(M_2)),$$

and set $\mathcal{K} := \ker \Psi_p \oplus \ker \tilde{\Psi}_p$.

Then, there exists a non-empty open subset $\mathcal{V} \subset M$ such that for every $p \in \mathcal{V}$,

$$\dim(\mathfrak{ho}_p^{\text{off}}(\nabla^C)) \geq r^\# := \text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}}) - \dim \mathcal{K}.$$

The bound depends only on $[\omega]_{\text{mixed}}$ and on the parallel-form strata $\mathcal{P}_k(M_i)$; hence it is invariant under any deformation of g_1 or g_2 that preserves these strata.

Proof. Let ω_h be the g -harmonic representative of $[\omega]$. Decompose ω_h according to Künneth and Hodge on the factors. For the $(2, 1)$ -part, write

$$\omega_h^{2,1} = \sum_{i=1}^{r_{2,1}} \alpha_i \wedge \beta_i, \quad (5.2)$$

with $\alpha_i \in \mathcal{H}^2(M_1)$ and $\beta_i \in \mathcal{H}^1(M_2)$ harmonic and chosen so that the number of terms is minimal. Similarly,

$$\omega_h^{1,2} = \sum_{j=1}^{r_{1,2}} \tilde{\alpha}_j \wedge \tilde{\beta}_j, \quad (5.3)$$

with $\tilde{\alpha}_j \in \mathcal{H}^1(M_1)$ and $\tilde{\beta}_j \in \mathcal{H}^2(M_2)$. The total number of simple mixed tensors equals $r := r_{2,1} + r_{1,2} = \text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}})$. Under the pure bidegree hypothesis, exactly one of $r_{2,1}$, $r_{1,2}$ is non-zero; the other vanishes together with the corresponding decomposition.

Set $r^\# := \text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}}) - \dim \mathcal{K}$. We will prove $\dim \mathfrak{ho}_p^{\text{off}}(\nabla^C) \geq r^\#$.

Consider the curvature formula (2.5) for $X \in V_1$ and $Y \in V_2$:

$$R^C(X, Y) = R^{LC}(X, Y) + \frac{1}{2}[(\nabla_X^{LC} T)(Y, \cdot, \cdot)^\# - (\nabla_Y^{LC} T)(X, \cdot, \cdot)^\#] + \frac{1}{4}Q_T(X, Y).$$

On a Riemannian product the mixed Riemann curvature vanishes: $R^{LC}(X, Y) = 0$ whenever $X \in V_1$ and $Y \in V_2$ (since $\nabla_X^{LC} Y = 0 = [X, Y]$ for vector fields tangent to different factors). For the remaining three terms we exploit the rigid pointwise algebraic structure that the surface hypothesis imposes on mixed 3-forms.

Since both factors are compact oriented surfaces, $\dim V_i = 2$, and the exterior power $\Lambda^2 V_1^*$ is one-dimensional, spanned by the area form vol_{V_1} . Every 3-form of bidegree $(2, 1)$ on M can therefore be written at each point p as

$$T_p^{2,1} = h(p) \text{vol}_{V_1} \wedge \tau_p, \quad (5.4)$$

where $h: M \rightarrow \mathbb{R}$ is a smooth function and τ is a smooth section of V_2^* (depending on the full point $p = (p_1, p_2)$). Likewise, every $(1, 2)$ -form satisfies $T_p^{1,2} = \tilde{h}(p) \tilde{\sigma}_p \wedge \text{vol}_{V_2}$ with $\tilde{\sigma} \in \Gamma(V_1^*)$. These representations hold for the full torsion T , not merely for its harmonic projection ω_h .

We denote by $J_i: V_i \rightarrow V_i$ the volume endomorphism of the oriented surface M_i , characterised by $g(J_i u, v) = \text{vol}_{V_i}(u, v)$; it is a skew-symmetric isometry with $J_i^2 = -\text{Id}$, and in particular $J_i X \perp X$ for every non-zero $X \in V_i$. Let $\pi_0: \mathfrak{so}(T_p M) \rightarrow \mathfrak{g}_0 := \mathfrak{so}(V_1) \oplus \mathfrak{so}(V_2)$ and $\pi_1: \mathfrak{so}(T_p M) \rightarrow \mathfrak{g}_1 := \mathfrak{so}^{\text{off}}(T_p M)$ denote the projections induced by the \mathbb{Z}_2 -grading of $\mathfrak{so}(T_p M)$.

A direct computation using (5.4) yields the following structure for the three surviving terms. The first, $\frac{1}{2}(\nabla_X^{LC} T)(Y, \cdot, \cdot)^\sharp$, is purely diagonal: applying the Leibniz rule to $\nabla_X^{LC}(h \text{vol}_{V_1} \wedge \tau)$ and using $\nabla_X^{LC} \text{vol}_{V_1} = 0$ (since $X \in V_1$ and the volume form of M_1 is parallel), one finds that this term is proportional to the volume endomorphism J_1 on V_1 (respectively J_2 on V_2 for the $(1, 2)$ -component of T); thus

$$\pi_0(R^C(X, Y)) = \lambda(X, Y) J_1 + \mu(X, Y) J_2 \quad (5.5)$$

for certain scalars λ, μ . The scalar λ arises from the $(2, 1)$ -component of T and its covariant derivatives, while μ arises from the $(1, 2)$ -component; under the pure bidegree hypothesis exactly one of λ, μ is identically zero. The second, $-\frac{1}{2}(\nabla_Y^{LC} T)(X, \cdot, \cdot)^\sharp$, is purely off-diagonal. Defining the 1-form $\rho_Y \in V_2^*$ by

$$\rho_Y := (Yh)\tau + h \nabla_Y^{g_2} \tau = \nabla_Y^{g_2}(h\tau), \quad (5.6)$$

the restriction to $\text{Hom}(V_1, V_2)$ reads

$$-\frac{1}{2}(\nabla_Y^{LC} T^{2,1})(X, Z, \cdot)^\sharp \Big|_{Z \in V_1} = -\frac{1}{2} g(J_1 X, Z) \rho_Y^\sharp \in V_2. \quad (5.7)$$

The third, $\frac{1}{4}Q_T(X, Y)$, is also purely off-diagonal (by a bidegree computation), with restriction

$$\frac{1}{4}Q_T(X, Y) \Big|_{Z \in V_1} = \frac{h^2}{4} \tau(Y) g(X, Z) \tau^\sharp \in V_2. \quad (5.8)$$

To derive (5.8), observe that the contortion tensor $K_X Y = \frac{1}{2}T(X, Y, \cdot)^\sharp$ acts on $Z \in V_1$ via the $(2, 1)$ -component as

$$K_Y^{(2,1)} Z = \frac{h}{2} \tau(Y) J_1 Z,$$

since $(\text{vol}_{V_1} \wedge \tau)(Y, Z, W) = \tau(Y) \text{vol}_{V_1}(Z, W)$ for $Y \in V_2$ and $Z, W \in V_1$. The iterated contortion $K_X(K_Y Z)$ with $X \in V_1$ then gives

$$K_X(K_Y Z) = \frac{h^2}{4} \tau(Y) g(X, Z) \tau^\sharp,$$

where we used $T^{(2,1)}(X, J_1 Z, W) = h \text{vol}_{V_1}(X, J_1 Z) \tau(W) = h g(J_1 X, J_1 Z) \tau(W) = h g(X, Z) \tau(W)$, the last equality following from J_1 being an isometry. Conversely, $K_Y(K_X Z) = 0$: since $K_X Z = \frac{h}{2} \text{vol}_{V_1}(X, Z) \tau^\sharp \in V_2$ and $T^{(2,1)}(Y, W, \cdot) = 0$ whenever $Y, W \in V_2$ (the form vol_{V_1} requires two arguments in V_1). Therefore $Q_T(X, Y)Z = K_X(K_Y Z) - K_Y(K_X Z)$ yields (5.8).

Combining (5.7) and (5.8), the full off-diagonal projection restricted to $\text{Hom}(V_1, V_2)$ is

$$\pi_1(R^C(X, Y))|_{V_1 \rightarrow V_2} : Z \mapsto -\frac{1}{2} g(J_1 X, Z) \rho_Y^\sharp + \frac{h^2}{4} \tau(Y) g(X, Z) \tau^\sharp. \quad (5.9)$$

The two summands act through the functionals $g(J_1 X, \cdot)$ and $g(X, \cdot)$ on V_1 respectively. Because $J_1 X \perp X$, these functionals are orthogonal, and the two contributions cannot cancel: if the map (5.9) vanishes for all $Z \in V_1$, then

$$\rho_Y^\sharp = 0 \quad \text{and} \quad h^2 \tau(Y) \tau^\sharp = 0 \quad (5.10)$$

must hold separately.

We now observe that the quadratic term (5.8) is non-vanishing on a non-empty open set. If $r_{2,1}^\sharp = 1$, the harmonic projection $\omega_h^{2,1} = c \cdot \text{vol}_{M_1} \wedge \beta$ is non-zero, so $[T^{2,1}] = [\omega_h^{2,1}] \neq 0$ in cohomology. Since $T^{2,1} = \omega_h^{2,1} + \eta^{2,1}$ with $\eta^{2,1}$ exact plus coexact, Hodge orthogonality gives $T^{2,1} \neq 0$ on M (otherwise $\omega_h^{2,1} = -\eta^{2,1}$ would be simultaneously harmonic and exact-plus-coexact, forcing $\omega_h^{2,1} = 0$). The set

$$\mathcal{W}^{2,1} := \{p \in M : T_p^{2,1} \neq 0\} \quad (5.11)$$

is therefore a non-empty open subset of M . On $\mathcal{W}^{2,1}$ we have $h(p) \neq 0$ and $\tau_p \neq 0$ (by (5.4)), so choosing $Y \in V_2$ with $\tau_p(Y) \neq 0$ yields $h^2 \tau(Y) \tau^\sharp \neq 0$; the second condition in (5.10) is therefore violated, and $\pi_1(R^C(X, Y)) \neq 0$. When T has pure bidegree $(1, 2)$, the symmetric argument for the $(1, 2)$ -component (using $J_2 Y \perp Y$) produces a non-empty open set $\mathcal{W}^{1,2}$ whenever $r_{1,2}^\sharp = 1$. Under the pure bidegree hypothesis, at most one of the sets $\mathcal{W}^{2,1}$, $\mathcal{W}^{1,2}$ is non-empty; both being simultaneously non-empty would require T to have mixed bidegree, contradicting the hypothesis.

In particular, the orthogonality of the two summands in (5.9) ensures that, on $\mathcal{W}^{2,1}$, the off-diagonal projection $\pi_1(R^C(X, Y))$ is non-zero regardless of the exact-plus-coexact part $\eta = T - \omega_h$. Indeed, h and τ in (5.4) are determined by the full torsion T , and on $\mathcal{W}^{2,1}$ one has $h(p) \neq 0$ and $\tau_p \neq 0$ by definition; the quadratic term $\frac{h^2}{4} \tau(Y) \tau^\sharp$ is therefore non-zero for a suitable choice of $Y \in V_2$, and acts through the functional $g(X, \cdot)$, which is orthogonal to $g(J_1 X, \cdot)$ appearing in the linear term. Thus, even if η alters the pointwise values of ρ_Y , h , and τ relative to those of the harmonic representative ω_h , the two contributions to the off-diagonal projection cannot cancel. In the admissible setting, $r^\sharp \leq 1$ per bidegree component, so this single protected direction suffices: the number of independent off-diagonal curvature directions is at least r^\sharp , computed from ω_h alone. The explicit construction of a *purely* off-diagonal holonomy element from the full curvature R^C is carried out in *Case 1* and *Case 2* below.

To organise these contributions, choose the harmonic forms $\{\alpha_i\}$ on M_1 and $\{\beta_i\}$ on M_2 mutually L^2 -orthonormal within their degrees, and likewise for $\{\tilde{\alpha}_j\}$, $\{\tilde{\beta}_j\}$. By minimality of the tensor-rank decomposition, the

sets $\{\alpha_i\}$ and $\{\beta_i\}$ are linearly independent as sections; hence there exists a dense open subset of points $p = (p_1, p_2) \in M_1 \times M_2$ where their evaluations $\{\alpha_i|_{p_1}\} \subset \Lambda^2 T_{p_1}^* M_1$ and $\{\beta_i|_{p_2}\} \subset T_{p_2}^* M_2$ are linearly independent. For any such p we pick vectors $X_i, Y_i \in T_{p_1} M_1$ and $Z_i \in T_{p_2} M_2$ satisfying $\alpha_i(X_i, Y_i)|_p = 1$ and $\beta_i(Z_i)|_p = 1$.

Define the linear map

$$\Psi_p : \mathcal{V}_{2,1} \longrightarrow \mathfrak{so}(T_p M), \quad \Psi_p \left(\sum_{i=1}^{r_{2,1}} c_i \alpha_i \otimes \beta_i \right) := \frac{1}{2} \sum_{i=1}^{r_{2,1}} c_i (\nabla^{LC} \omega_h)^{\text{off}}(X_i, Y_i, Z_i), \quad (5.12)$$

where the vectors $X_i, Y_i \in T_{p_1} M_1$ and $Z_i \in T_{p_2} M_2$ satisfy $\alpha_i(X_i, Y_i) = 1$ and $\beta_i(Z_i) = 1$. The notation on the right-hand side denotes the off-diagonal component of the curvature contribution from the $\nabla^{LC} \omega_h$ -term in (2.5), and its precise meaning is established through the kernel characterization in *Claim 5.2.1*.

Claim 5.2.1 (Independence of choices). The subspace $\ker \Psi_p \subset \mathcal{V}_{2,1}$ is independent of the choice of normalizing vectors $\{X_i, Y_i, Z_i\}$ in (5.12).

Proof of Claim. Writing $\omega_h^{2,1} = \sum_{i=1}^{r_{2,1}} \alpha_i \wedge \beta_i$, the product rule gives

$$(\nabla^{LC} \omega_h)^{\text{off}} = \sum_{i=1}^{r_{2,1}} (\nabla^{g_1} \alpha_i) \wedge \beta_i + \alpha_i \wedge (\nabla^{g_2} \beta_i).$$

For vectors $X_i, Y_i \in T_{p_1} M_1$, $Z_i \in T_{p_2} M_2$ with $\alpha_i(X_i, Y_i) = 1$, $\beta_i(Z_i) = 1$, the contraction $\iota_{Z_i} \iota_{Y_i} \iota_{X_i}$ extracts precisely the i -th component:

$$(\nabla^{LC} \omega_h)^{\text{off}}(X_i, Y_i; Z_i) = (\nabla^{g_1} \alpha_i) \wedge \beta_i + \alpha_i \wedge (\nabla^{g_2} \beta_i).$$

Summing over i with coefficients c_i yields that $\Psi_p(\xi) = 0$ is equivalent to the tensor equation

$$\sum_{i=1}^{r_{2,1}} c_i \nabla^{LC} (\alpha_i \wedge \beta_i) = 0, \quad (5.13)$$

as we now verify in the admissible setting. Since $r_{2,1} = 1$, there is a single term $\alpha_1 \wedge \beta_1$. Because $\alpha_1 \in H^2(M_1) = \mathbb{R} \cdot \text{vol}_{M_1}$ is parallel, the Leibniz rule gives $\nabla^{LC} (\alpha_1 \wedge \beta_1) = \alpha_1 \wedge (\nabla^{g_2} \beta_1)$. The evaluation on the V_1 -slots via any pair (X, Y) with $\alpha_1(X, Y) = 1$ yields the fixed factor $\alpha_1(X, Y) = 1$, so the dependence on the choice of normalising vectors reduces entirely to the V_2 -slot. Different choices $Z, Z' \in V_2$ with $\beta_1(Z) = \beta_1(Z') = 1$ differ by an element $v \in \ker \beta_1$, and the evaluation changes by the additive term $(\nabla^{g_2} \beta_1)(v, \cdot)$. Now, since β_1 is harmonic on the compact oriented surface M_2 , the conditions $d\beta_1 = 0$ and $\delta\beta_1 = 0$ imply that $\nabla^{g_2} \beta_1$ is a symmetric trace-free $(0, 2)$ -tensor. In dimension 2, such a tensor has rank 0 or rank 2: in any orthonormal frame its matrix takes the form $\begin{pmatrix} a & b \\ b & -a \end{pmatrix}$ with determinant $-(a^2 + b^2)$, which vanishes if and only if the tensor itself vanishes. Consequently, if $\nabla^{g_2} \beta_1|_{p_2} \neq 0$ then $(\nabla^{g_2} \beta_1)(Z, \cdot) \neq 0$ for every non-zero Z , regardless of

the choice; if $\nabla^{g_2}\beta_1|_{p_2} = 0$ then $(\nabla^{g_2}\beta_1)(Z, \cdot) = 0$ for every Z . In either case the condition $\Psi_p(\xi) = 0$ is independent of the normalizing vectors and equivalent to the vanishing of (5.13), which is manifestly independent of the vectors $\{X_i, Y_i, Z_i\}$. Hence $\ker \Psi_p$ is well-defined. \square

Now, although the right-hand side of (5.12) depends on the choice of vectors, by *Claim 5.2.1* the condition $\Psi_p(\xi) = 0$ is independent of that choice.

If $\xi \in \ker \Psi_p$, then by definition

$$\frac{1}{2} \sum_i c_i (\nabla^{LC} \omega_h)^{\text{off}}(X_i, Y_i, Z_i) = 0.$$

Separating the bidegrees in (5.13) yields the two independent conditions

$$\sum_{i=1}^{r_{2,1}} c_i (\nabla^{g_1} \alpha_i) \wedge \beta_i = 0, \quad \sum_{i=1}^{r_{2,1}} c_i \alpha_i \wedge (\nabla^{g_2} \beta_i) = 0. \quad (5.14)$$

In the admissible setting $r_{2,1} \leq 1$ (Definition 2.2 and Lemma 2.3), we establish the forward inclusion

$$\ker \Psi_p \subseteq \mathcal{V}_{2,1} \cap (\mathcal{P}_2(M_1) \otimes \mathcal{P}_1(M_2))$$

on a non-empty open subset of M , by a direct pointwise evaluation of (5.14). If $r_{2,1} = 0$ the inclusion is vacuous (and holds on all of M). If $r_{2,1} = 1$, write $\xi = c_1(\alpha_1 \otimes \beta_1)$ and suppose $\xi \in \ker \Psi_p$. We distinguish two cases according to a global dichotomy: either α_1 is parallel on all of M_1 , or it is not.

- (i) If $\alpha_1 \notin \mathcal{P}_2(M_1)$, then $\nabla^{g_1} \alpha_1 \neq 0$, so by Lemma 5.1 the set $\{p_1 \in M_1 : (\nabla^{g_1} \alpha_1)(p_1) \neq 0\}$ is a non-empty open subset of M_1 . Since $\beta_1 \neq 0$, the set $\{p_2 \in M_2 : \beta_1(p_2) \neq 0\}$ is likewise non-empty and open. Define

$$\mathcal{O}' := \{(p_1, p_2) \in M : (\nabla^{g_1} \alpha_1)(p_1) \neq 0 \text{ and } \beta_1(p_2) \neq 0\},$$

which is a non-empty open subset of M . For any $p \in \mathcal{O}'$, if $\xi \in \ker \Psi_p$, then (5.14) evaluated at this same p gives

$$c_1 (\nabla^{g_1} \alpha_1)(p_1) \wedge \beta_1(p_2) = 0.$$

Both factors are non-zero by definition of \mathcal{O}' , so $c_1 = 0$, hence $\xi = 0 \in \mathcal{P}_2(M_1) \otimes \mathcal{P}_1(M_2)$.

- (ii) If $\alpha_1 \in \mathcal{P}_2(M_1)$, the same argument is applied to β_1 via the second equation in (5.14). If $\beta_1 \notin \mathcal{P}_1(M_2)$, define

$$\mathcal{O}'' := \{(p_1, p_2) \in M : \alpha_1(p_1) \neq 0 \text{ and } (\nabla^{g_2} \beta_1)(p_2) \neq 0\},$$

which is a non-empty open subset of M (since $\alpha_1 \in \mathcal{P}_2(M_1)$ is a non-zero parallel 2-form on a surface, it never vanishes, and Lemma 5.1 gives the non-emptiness of the second factor). For any $p \in \mathcal{O}''$, if $\xi \in \ker \Psi_p$, the second equation in (5.14) at this same p forces $c_1 = 0$, hence $\xi = 0 \in \mathcal{P}_2(M_1) \otimes \mathcal{P}_1(M_2)$. If $\beta_1 \in \mathcal{P}_1(M_2)$, then

$\xi = c_1(\alpha_1 \otimes \beta_1) \in \mathcal{P}_2(M_1) \otimes \mathcal{P}_1(M_2)$ for every c_1 , and the inclusion holds on all of M .

In all cases there exists a non-empty open subset of M (namely \mathcal{O}' , \mathcal{O}'' , or all of M respectively) on which $\ker \Psi_p \subseteq \mathcal{V}_{2,1} \cap (\mathcal{P}_2(M_1) \otimes \mathcal{P}_1(M_2))$.

For the opposite inclusion, if $\xi \in \mathcal{V}_{2,1} \cap (\mathcal{P}_2(M_1) \otimes \mathcal{P}_1(M_2))$, then ξ can be written as a finite sum $\sum_k A_k \otimes B_k$ with $A_k \in \mathcal{P}_2(M_1)$, $B_k \in \mathcal{P}_1(M_2)$. For each term, $\nabla^{LC}(A_k \wedge B_k) = 0$ because A_k and B_k are parallel, hence $\Psi_p(A_k \otimes B_k) = 0$. By linearity, $\Psi_p(\xi) = 0$, so $\xi \in \ker \Psi_p$.

Therefore, on the non-empty open set \mathcal{O}' (resp. \mathcal{O}'' , resp. all of M when both forms are parallel),

$$\ker \Psi_p = \mathcal{V}_{2,1} \cap (\mathcal{P}_2(M_1) \otimes \mathcal{P}_1(M_2)).$$

An identical argument applied to the (1,2)-component yields, on a corresponding non-empty open subset,

$$\ker \tilde{\Psi}_p = \mathcal{V}_{1,2} \cap (\mathcal{P}_1(M_1) \otimes \mathcal{P}_2(M_2)).$$

Now consider the induced map on the quotient

$$\bar{\Psi}_p : \mathcal{V}_{2,1} / \ker \Psi_p \longrightarrow \mathfrak{so}(T_p M), \quad \bar{\Psi}_p([\xi]) := \Psi_p(\xi).$$

The map $\bar{\Psi}_p$ is well-defined and linear. It is injective because $\bar{\Psi}_p([\xi]) = 0$ implies $\Psi_p(\xi) = 0$, hence $\xi \in \ker \Psi_p$ and $[\xi] = 0$. Consequently, $\bar{\Psi}_p$ sends a basis of the quotient to a set of linearly independent off-diagonal curvature operators, and we obtain

$$\dim(\text{im } \bar{\Psi}_p) = \dim(\mathcal{V}_{2,1} / \ker \Psi_p) = r_{2,1} - \dim \ker \Psi_p. \quad (5.15)$$

An entirely analogous argument for the (1,2)-component gives

$$\dim(\text{im } \tilde{\bar{\Psi}}_p) = r_{1,2} - \dim \ker \tilde{\Psi}_p. \quad (5.16)$$

Let $U = \text{span}\{\alpha_i\} \subset \mathcal{H}^2(M_1)$ and $V = \text{span}\{\beta_i\} \subset \mathcal{H}^1(M_2)$. Define

$$U_{\parallel} = U \cap \mathcal{P}_2(M_1), \quad V_{\parallel} = V \cap \mathcal{P}_1(M_2),$$

and set $s = \dim U_{\parallel}$, $t = \dim V_{\parallel}$. Choose any basis $\{\alpha_1, \dots, \alpha_s\}$ of U_{\parallel} and complete it to a basis of U by adding vectors $\{\alpha_{s+1}, \dots, \alpha_{r_{2,1}}\}$ whose classes form a basis of U/U_{\parallel} (hence linearly independent modulo $\mathcal{P}_2(M_1)$).

By Lemma 5.1 applied to the subspace $U \subset \mathcal{H}^2(M_1)$ with parallel part of dimension s , there exists a non-empty open set $\mathcal{U}_1^{2,1} \subset M_1$ such that for every $p_1 \in \mathcal{U}_1^{2,1}$ the covariant derivatives $\nabla^{g_1} \alpha_i(p_1)$ (for $i = s+1, \dots, r_{2,1}$) are linearly independent.

Similarly, applying Lemma 5.1 to $\text{span}\{\beta_i\} \subset \mathcal{H}^1(M_2)$ with parallel part of dimension t yields a non-empty open set $\mathcal{U}_2^{2,1} \subset M_2$ such that for every $p_2 \in \mathcal{U}_2^{2,1}$ the covariant derivatives $\nabla^{g_2} \beta_i(p_2)$ (for $i = t+1, \dots, r_{2,1}$) are linearly independent.

Under the pure bidegree hypothesis, only one family of harmonic forms is present. If $T = T^{2,1}$, the forms $\{\tilde{\alpha}_j\}$, $\{\tilde{\beta}_j\}$ are absent, and we set

$$\mathcal{U}_1 := \mathcal{U}_1^{(2,1)} \subset M_1, \quad \mathcal{U}_2 := \mathcal{U}_2^{(2,1)} \subset M_2;$$

if $T = T^{1,2}$, the forms $\{\alpha_i\}$, $\{\beta_i\}$ are absent, and we set

$$\mathcal{U}_1 := \mathcal{U}_1^{(1,2)} \subset M_1, \quad \mathcal{U}_2 := \mathcal{U}_2^{(1,2)} \subset M_2.$$

In either case, \mathcal{U}_1 and \mathcal{U}_2 are non-empty open subsets of M_1 and M_2 respectively (by Lemma 5.1), and the product $\mathcal{U} := \mathcal{U}_1 \times \mathcal{U}_2$ is a non-empty open subset of $M = M_1 \times M_2$.

Since $\mathcal{O}' \cap \mathcal{U} = \mathcal{U}_1 \times (\{p_2 \in M_2 : \beta_1(p_2) \neq 0\} \cap \mathcal{U}_2)$, it suffices to show $\{\beta_1 \neq 0\} \cap \mathcal{U}_2 \neq \emptyset$. Were $\beta_1 \equiv 0$ on \mathcal{U}_2 , then $\nabla^{g_2} \beta_1 \equiv 0$ on \mathcal{U}_2 (since β_1 vanishes on an open set), contradicting the definition of \mathcal{U}_2 . Hence $\mathcal{O} := \mathcal{O}' \cap \mathcal{U}$ is a non-empty open subset of M . (In the case $\alpha_1 \in \mathcal{P}_2(M_1)$ and $\beta_1 \notin \mathcal{P}_1(M_2)$, one has $\mathcal{O}'' = M_1 \times \mathcal{U}_2 \supseteq \mathcal{U}$, so $\mathcal{O}'' \cap \mathcal{U} = \mathcal{U}$; when both forms are parallel, the kernel identification holds on all of M and $\mathcal{O} = \mathcal{U}$.)

For any point $p = (p_1, p_2) \in \mathcal{O}$ the kernel identification and the linear independence of covariant derivatives hold simultaneously; consequently the linear maps Ψ_p and $\tilde{\Psi}_p$ satisfy

$$\dim \operatorname{im} \Psi_p = r_{2,1} - \dim \ker \Psi_p, \quad \dim \operatorname{im} \tilde{\Psi}_p = r_{1,2} - \dim \ker \tilde{\Psi}_p,$$

with $\ker \Psi_p = \mathcal{V}_{2,1} \cap (\mathcal{P}_2(M_1) \otimes \mathcal{P}_1(M_2))$.

Hence, at every $p \in \mathcal{O}$, the off-diagonal part of the linear term $\frac{1}{2}(\nabla^{LC}T)^{\text{off}}$ generates at least

$$(r_{2,1} - \dim \ker \Psi_p) + (r_{1,2} - \dim \ker \tilde{\Psi}_p) = r - \dim \mathcal{K}$$

linearly independent curvature operators (each having non-trivial off-diagonal projection). Under the pure bidegree hypothesis exactly one of the two summands is non-zero; the other vanishes identically.

It remains to produce a purely off-diagonal element of the holonomy algebra. Under the pure bidegree hypothesis exactly one of $r_{2,1}$, $r_{1,2}$ is non-zero, and for admissible manifolds each is at most 1; consequently $r^\# = r - \dim \mathcal{K} \leq 1$, and it suffices to exhibit a single purely off-diagonal holonomy element on a non-empty open set.

Without loss of generality assume T has pure bidegree $(2,1)$, so that $r_{2,1} \geq 1$ and $r_{1,2} = 0$ (the case of pure bidegree $(1,2)$ is entirely symmetric). We work on the non-empty open set $V := W^{2,1}$, where $W^{2,1}$ is the non-empty open set introduced in 5.11. At every point $p \in V$ the torsion satisfies $h(p) \neq 0$ and $\tau_p \neq 0$ by definition of $W^{2,1}$.

Consider the mixed curvature span

$$\mathcal{S}_p := \operatorname{span}\{R^C(X, Y) : X \in V_1, Y \in V_2\}.$$

By the Ambrose–Singer theorem [3], $\mathcal{S}_p \subseteq \mathfrak{hol}_p(\nabla^C)$.

We now show that $\mathfrak{hol}_p(\nabla^C)$ contains at least one purely off-diagonal element on a non-empty open subset of M .

Suppose $r_{2,1}^\# = 1$. Recall that at every point $p \in \mathcal{W}^{2,1}$ we have $h(p) \neq 0$ and $\tau_p \neq 0$. We distinguish two cases according to whether p lies in the interior of $\{T^{1,2} = 0\}$.

Case 1: $p \in \mathcal{W}^{2,1} \setminus \overline{\mathcal{W}^{1,2}}$. Since $\mathcal{W}^{1,2} = \{q \in M : T_q^{1,2} \neq 0\}$ is open, its complement $\{T^{1,2} = 0\}$ is closed and its interior is $M \setminus \overline{\mathcal{W}^{1,2}}$. At every point $p \in \mathcal{W}^{2,1} \setminus \overline{\mathcal{W}^{1,2}}$ the $(1,2)$ -component $T^{1,2}$ vanishes identically on an open neighbourhood of p ; consequently $\nabla^{\text{LC}} T^{1,2}|_p = 0$ as well (the derivative of a section that vanishes on an open set is zero on that set).

It follows that both the quadratic torsion terms involving $T^{1,2}$ and the $\nabla^{\text{LC}} T^{1,2}$ -terms in the curvature formula (2.5) vanish at p . The diagonal projection therefore reduces to

$$\pi_0(R^C(X, Y)) = \lambda(X, Y) J_1$$

(the $\mu(X, Y) J_2$ contribution is absent because it arises entirely from $T^{1,2}$ and its covariant derivatives, both of which vanish at p). Choose $Y \in V_2$ with $\tau_p(Y) \neq 0$ (possible because $p \in \mathcal{W}^{2,1}$ implies $\tau_p \neq 0$). The linear functional $\lambda(\cdot, Y) : V_1 \rightarrow \mathbb{R}$ has kernel of dimension at least 1 (since $\dim V_1 = 2$). Pick $X_0 \in \ker(\lambda(\cdot, Y)) \setminus \{0\}$. Then:

$$\text{Diagonal part: } \pi_0(R^C(X_0, Y)) = \lambda(X_0, Y) J_1 = 0.$$

Off-diagonal part evaluated on $Z = X_0$: by (5.9), the two contributions are $-\frac{1}{2} g(J_1 X_0, X_0) \rho_Y^\sharp$ and $\frac{h^2}{4} \tau(Y) g(X_0, X_0) \tau^\sharp$. The first vanishes because $g(J_1 X_0, X_0) = 0$ (skew-symmetry of J_1); the second equals $\frac{h^2}{4} \tau(Y) |X_0|^2 \tau^\sharp \neq 0$ (since $h \neq 0$, $\tau(Y) \neq 0$, and $|X_0|^2 > 0$).

Therefore $R^C(X_0, Y)$ is a non-zero element of $\mathfrak{so}^{\text{off}}(T_p M)$. By the Ambrose–Singer theorem [3], $R^C(X_0, Y) \in \mathfrak{hol}_p(\nabla^C)$. Since $\mathfrak{hol}_p^{\text{off}}(\nabla^C) = \mathfrak{hol}_p(\nabla^C) \cap \mathfrak{so}^{\text{off}}(T_p M)$ (Definition 3.2), we conclude that $R^C(X_0, Y) \in \mathfrak{hol}_p^{\text{off}}(\nabla^C)$, whence $\dim \mathfrak{hol}_p^{\text{off}}(\nabla^C) \geq 1$.

Case 2: $p \in \mathcal{W}^{2,1} \cap \mathcal{W}^{1,2}$ (both components present). This case is vacuous under the pure bidegree hypothesis. We show that the hypothesis $T \in \Gamma(\Lambda^2 V_1^* \otimes V_2^*)$ (the argument for $T \in \Gamma(V_1^* \otimes \Lambda^2 V_2^*)$ is entirely symmetric) eliminates the overlap $\mathcal{W}^{2,1} \cap \mathcal{W}^{1,2}$ altogether.

The pure bidegree condition requires that for every $p \in M$ and all $X \in V_1(p)$, $Y_1, Y_2 \in V_2(p)$ one has $T_p(X, Y_1, Y_2) = 0$, which is precisely the statement $T^{1,2} \equiv 0$ as a section of $V_1^* \otimes \Lambda^2 V_2^*$. Consequently

$$\mathcal{W}^{1,2} = \{p \in M : T^{1,2}(p) \neq 0\} = \emptyset,$$

and therefore $\mathcal{W}^{2,1} \cap \mathcal{W}^{1,2} = \emptyset$. This is not a generic or measure-theoretic statement: the overlap is empty at every point of M . In particular, every $p \in \mathcal{W}^{2,1}$ admits a neighbourhood on which $T^{1,2}$ vanishes identically, so the hypotheses of Case 1 are satisfied at all such points.

Conversely, if $\mathcal{W}^{2,1} \cap \mathcal{W}^{1,2} \neq \emptyset$ then there exists a point at which both $T^{2,1}$ and $T^{1,2}$ are non-zero, whence T cannot have pure bidegree. The pure bidegree hypothesis is therefore the precise condition that renders Case 2 vacuous.

Since $T = T^{2,1}$ is a continuous section of $\Lambda^2 V_1^* \otimes V_2^*$, the set

$$\mathcal{W}^{2,1} = \{p \in M : T^{2,1}(p) \neq 0\}$$

is open in M . We claim that it is non-empty. Because $T = T^{2,1}$, the de Rham class $[T^\flat]$ lies in $H^2(M_1) \otimes H^1(M_2)$ and coincides with the harmonic representative $[\omega_h]$, which is a non-trivial mixed class by hypothesis. Were $T^{2,1}$ identically zero, then $T^\flat = 0$ and $[T^\flat] = 0$, contradicting the non-triviality of $[\omega_h]$. Hence $\mathcal{W}^{2,1} \neq \emptyset$.

Set $\mathcal{V} := \mathcal{W}^{2,1}$ (in the symmetric case $T = T^{1,2}$ one sets $V := W^{1,2}$). This is a non-empty open subset of M , and at every point of \mathcal{V} the hypotheses of *Case 1* are in force (since $\mathcal{W}^{1,2} = \emptyset$). Applying the conclusion of *Case 1* we obtain, for every $p \in \mathcal{V}$,

$$\mathfrak{hol}_p^{\text{off}}(\nabla^C) \neq \{0\},$$

together with the rank estimate

$$\dim(\mathfrak{hol}_p^{\text{off}}(\nabla^C)) \geq r^\sharp := \text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}}) - \dim \mathcal{K}.$$

The quantity r^\sharp depends only on the mixed Künneth class $[\omega]_{\text{mixed}}$ and on the parallel-form strata $\mathcal{P}_k(M_i)$; it is therefore invariant under any deformation of g_1 or g_2 that preserves these strata. □

Corollary 5.3. *Under the hypotheses of Theorem 5.2, if $\text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}}) \geq 1$, then $\dim \mathfrak{hol}_p^{\text{off}}(\nabla^C) \geq 1$ on the non-empty open subset $\mathcal{V} \subset M$ of Theorem 5.2. In particular, the holonomy of ∇^C is irreducible with respect to the product splitting $V_1 \oplus V_2$.*

Proof. Immediate from *Case 1* of the proof of Theorem 5.2: the quadratic torsion term $\frac{1}{4}Q_T$ produces a non-zero purely off-diagonal element of $\mathfrak{hol}_p(\nabla^C)$ at every point of $\mathcal{W}^{2,1} \neq \emptyset$ (respectively $\mathcal{W}^{1,2} \neq \emptyset$), using only the non-vanishing of $T^{2,1}$ (respectively $T^{1,2}$), which is guaranteed by $\text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}}) \geq 1$. □

Remark 5.4 (Role of the harmonic projection). The invariant r^\sharp is defined through the harmonic projection ω_h alone, not through the full torsion $T = \omega_h + d\alpha + \delta\beta$. This is essential for invariance: for a fixed class $[\omega]$, infinitely many torsion forms represent $[\omega]$ (corresponding to different exact and coexact parts), so any quantity depending on the full torsion would not be an invariant of the cohomology class.

As a consequence, r^\sharp captures only the off-diagonal holonomy forced by the cohomological datum. More precisely, the linear map Ψ_p defined in 5.12 serves to *count* the number of independent off-diagonal directions: its injectivity on the quotient $\mathcal{V}_{2,1}/\ker \Psi_p$, which holds on a non-empty open subset of M (see the proof of Theorem 5.2), yields $\dim \text{im } \Psi_p = r_{2,1} - \dim \ker \Psi_p$. The *existence* of a purely off-diagonal holonomy element is then established by *Case 1* of the proof through the quadratic torsion term $\frac{1}{4}Q_T$, which depends on the pointwise values of T rather than on its cohomology

class alone. Since Q_T may generate off-diagonal curvature beyond what is detected by r^\sharp , the invariant r^\sharp provides a lower bound, not an exact count, for the off-diagonal holonomy dimension. More precisely, r^\sharp certifies irreducibility that arises when the irreducible de Rham geometry of a factor prevents the cohomological coupling from being absorbed by the parallel-form strata; when $r^\sharp = 0$, no such obstruction to absorption is present, and the irreducibility, though still present via the quadratic torsion term Q_T (as discussed above), does not arise from this mechanism (see Remark 5.6). In this sense, r^\sharp quantifies the *transverse irreducibility of the connection* (Remark 5.6).

However, in the generic situation where $\nabla^{LC}T \neq 0$, the curvature operator $R^C(X, Y)$ has both diagonal and off-diagonal components, and the off-diagonal projection of a holonomy element is not itself a holonomy element in general; extracting a *purely* off-diagonal element from the full curvature requires the algebraic mechanism of *Theorem 5.2*.

Remark 5.5 (r^\sharp as a property of the pair (M, g)). The condition $r^\sharp([\omega], g) \geq 1$ for some class $[\omega] \in H^3(M)$ is a property of the pair (M, g) in admissible manifolds. When satisfied, r^\sharp provides a non-trivial lower bound $\dim \mathfrak{ho}_p^{\text{off}} \geq 1$, certifying irremovable irreducibility for every metric connection with totally antisymmetric torsion of pure bidegree representing $[\omega]$. When $r^\sharp([\omega], g) = 0$ for every class $[\omega]$, no such assurance is provided: for instance, on $S^2 \times S^2$ the group $H^3(M) = 0$, so no metric connection with totally antisymmetric torsion can represent a non-trivial cohomology class.

Remark 5.6 (Geometric interpretation of r^\sharp and the obstruction space \mathcal{K}). The invariant $r^\sharp = ([\omega]_{\text{mixed}}) - \dim \mathcal{K}$ is computable from the cohomology class $[\omega_h]$ and the parallel-form strata $\mathcal{P}_k(M_i)$ alone, without knowledge of the specific torsion representative T . The correction $-\dim \mathcal{K}$ admits a natural geometric interpretation in terms of the de Rham decomposition of the factors.

When $\dim \mathcal{K} > 0$, the subspace \mathcal{K} contains at least one simple tensor $\alpha_k \otimes \beta_k$ with $\alpha_k \in \mathcal{P}_2(M_i)$ and $\beta_k \in \mathcal{P}_1(M_j)$, i.e. both factors are parallel with respect to the Levi-Civita connections of the factor metrics. In the admissible setting, the condition on α_k is automatic ($\mathcal{P}_2(M_i) = \mathbb{R} \cdot \text{vol}_{M_i}$ for every compact oriented surface), and the operative restriction concerns only β_k . A non-zero parallel 1-form β_k on a compact Riemannian manifold defines a parallel vector field β_k^\sharp , which by the de Rham decomposition theorem determines a flat sub-factor in the universal cover: \widetilde{M}_j splits isometrically as $\widetilde{M}_j \cong \widetilde{N}_j \times \mathbb{R}_{\beta_k}$. In the admissible setting, $\mathcal{P}_1(M_j) \neq \{0\}$ forces M_j to be a flat torus, and the splitting descends to the manifold itself: $M_j \cong N_j \times S_{\beta_k}^1$. The component $\alpha_k \otimes \beta_k$ of the torsion class then couples M_i exclusively with the flat direction $S_{\beta_k}^1$, without involving N_j .

The components not in \mathcal{K} (those with $\beta_l \notin \mathcal{P}_1(M_j)$, or $\alpha_l \notin \mathcal{P}_2(M_i)$ in higher-dimensional settings) couple M_i with the de Rham irreducible part of M_j , the geometrically non-trivial part of the factor.

The invariant r^\sharp therefore measures the dimension of the part of the mixed cohomological coupling that is *transverse* to the de Rham splitting, i.e. the part that involves the irreducible factor. The correction $\dim \mathcal{K}$ measures the dimension of the part *aligned* with the flat sub-factors.

This interpretation clarifies the meaning of the two regimes:

- (i) When $r^\sharp \geq 1$: the irreducible de Rham geometry of a factor prevents the cohomological coupling from being absorbed by the parallel-form strata. Theorem 5.2 certifies that this non-absorption forces irreducibility of the holonomy, which cannot be removed by metric deformations preserving the parallel-form strata. We refer to this as *transverse irreducibility of the connection*.
- (ii) When $r^\sharp = 0$: no irreducible de Rham factor prevents absorption; the entire cohomological coupling is absorbed by the parallel-form strata. Transverse irreducibility is absent. This does not imply that the holonomy is reducible, the quadratic torsion term $\frac{1}{4}Q_T$ can still generate off-diagonal curvature, as observed in Remark 5.4, but signals that the irreducibility does not arise from the mechanism detected by r^\sharp .

In the admissible setting (products of surfaces), the distinction reduces to the dichotomy of Remark 5.10. For higher-dimensional factors, the de Rham decomposition is richer: k independent parallel 1-forms on M_j induce a splitting of the universal cover $\widetilde{M}_j \cong \widetilde{N}_j \times \mathbb{R}^k$; when this splitting descends to the manifold (e.g. under a simple connectivity hypothesis or for factors with special holonomy), one obtains $M_j \cong N_j \times T^k$, and $\dim \mathcal{K}$ measures the dimension of the part of the mixed class aligned with T^k . In such settings, the invariant r^\sharp acquires a structural meaning that goes beyond the bound itself: it detects how deeply the torsion coupling penetrates the de Rham irreducible geometry of the factors.

Remark 5.7 (Geometric Significance of Invariance). The invariance of the lower bound r^\sharp is a non-trivial property that holds under the Assumption 2.2. In those cases, the bound depends only on:

- (i) the mixed class $[\omega]_{\text{mixed}}$,
- (ii) the parallel-form strata $\mathcal{P}_k(M_i)$ of the factor metrics.

Consequently, r^\sharp is invariant under any deformation of g_1 or g_2 that preserves these strata. This stability is relevant for manifolds with special holonomy where $\mathcal{P}_k(M_i)$ are non-trivial.

Remark 5.8 (Forced irreducibility and quantified off-diagonal holonomy). Since ∇^C is metric, any non-trivial off-diagonal holonomy implies that $\text{hol}_p(\nabla^C)$ acts irreducibly on $T_p M$. Theorem 5.2 quantifies this “forced irreducibility”

principle in terms of the mixed tensor rank, modulo the obstruction space

$$\mathcal{K} = \ker \Psi_p \oplus \ker \tilde{\Psi}_p,$$

where, in the notation of the theorem,

$$\ker \Psi_p = \mathcal{V}_{2,1} \cap (\mathcal{P}_2(M_1) \otimes \mathcal{P}_1(M_2)), \quad \ker \tilde{\Psi}_p = \mathcal{V}_{1,2} \cap (\mathcal{P}_1(M_1) \otimes \mathcal{P}_2(M_2)).$$

Specifically, the dimension of the off-diagonal holonomy subspace satisfies

$$\dim \mathfrak{hol}_p^{\text{off}}(\nabla^C) \geq r^\sharp := \text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}}) - \dim \mathcal{K},$$

and the bound certifies irreducibility when $r^\sharp > 0$, i.e. when $\text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}}) > \dim \mathcal{K}$; when $r^\sharp = 0$ the inequality is trivially satisfied; the geometric content of this regime is discussed in Remark 5.6.

In other words, a non-zero value of the corrected rank r^\sharp guarantees that the holonomy representation cannot split according to the product decomposition $T_p M = V_1 \oplus V_2$; the connection is geometrically “entangled” across the two factors.

Remark 5.9 (Analytic approximation as a tool of proof). The proof of the density statement in Lemma 5.1 makes essential use of real-analytic approximations of the smooth metric. Real-analytic metrics are dense in the C^∞ topology (Greene–Jacobowitz [5]), and the relevant quantities – the dimension of the space of parallel forms, the harmonic representatives of fixed cohomology classes, and the covariant derivatives of those harmonic forms – depend continuously on the metric in the appropriate topologies. The non-emptiness conclusion of Lemma 5.1, which is all that is required in the proof of Theorem 5.2, holds for arbitrary smooth metrics without recourse to analytic approximation.

In Lemma 5.1 the reduction to an analytic metric allows us to apply the identity principle for analytic functions: if an analytic section of a vector bundle vanishes on an open set, it vanishes identically. This principle is used to show that a relation among the covariant derivatives of harmonic forms on an open set forces the existence of a parallel combination of the forms, contradicting the hypothesis of linear independence modulo parallel forms. For a general smooth metric, the non-emptiness of \mathcal{U} is immediate from the non-parallelism hypothesis; the density conclusion is established only under the additional assumption that g_1 is real-analytic.

In Theorem 5.2, only the non-emptiness of \mathcal{U} is needed: in the admissible setting the condition $r - s \leq 1$ ensures that a single evaluation point suffices for the kernel identification, and the purely off-diagonal holonomy element is constructed on $W^{2,1}$ without reference to \mathcal{U} .

Thus analyticity serves as a tool to establish the stronger density statement in Lemma 5.1; the final statements of Theorem 5.2 are valid for arbitrary smooth product metrics and do not require any analyticity assumption on the metric or the torsion.

Remark 5.10 (Dichotomy of the obstruction). For admissible product manifolds, the obstruction space \mathcal{K} is controlled by a single topological datum:

the vanishing or non-vanishing of $\mathcal{P}_1(M_i)$. Since $\mathcal{P}_2(M_i) = \mathbb{R} \cdot \text{vol}_{M_i}$ for every compact oriented surface, the obstruction kernels reduce to

$$\ker \Psi_p \neq \{0\} \iff \beta \in \mathcal{P}_1(M_2), \quad \ker \tilde{\Psi}_p \neq \{0\} \iff \alpha \in \mathcal{P}_1(M_1),$$

where α, β are the harmonic forms appearing in the minimal decomposition of $[\omega]_{\text{mixed}}$. A compact oriented surface F satisfies $\mathcal{P}_1(F) = \{0\}$ if and only if one of the following holds: $g(F) = 0$ (since $H^1(S^2) = \{0\}$); $g(F) \geq 2$ (by the Poincaré–Hopf theorem: $\chi(F) \neq 0$ forbids nowhere-vanishing vector fields, hence parallel 1-forms); or $g(F) = 1$ and the metric is not flat (a non-zero parallel 1-form on a surface forces the Gaussian curvature to vanish identically). In the context of Theorem 5.2, the case $g(F) = 0$ is vacuous: a non-trivial mixed class in $H^2(M_1) \otimes H^1(M_2)$ requires $H^1(M_2) \neq \{0\}$, hence $g(M_2) \geq 1$.

This yields a sharp dichotomy. When at least one factor has genus $g \geq 2$, the obstruction space \mathcal{K} is at most one-dimensional regardless of the metric, and $r^\sharp = 1$ whenever the mixed class $[\omega]_{\text{mixed}}$ is non-trivial in the corresponding Künneth component; in particular, on $\Sigma_{g_1} \times \Sigma_{g_2}$ with $g_1, g_2 \geq 2$, the bound $r^\sharp = 1$ is *purely topological* and forces irreducibility for every cohomologically calibrated connection with torsion representing a non-trivial mixed class, independently of the choice of product metric. The only setting where the metric controls the transition $r^\sharp = 0 \leftrightarrow r^\sharp = 1$ is when the relevant factor is a torus T^2 : the parallel-form stratum $\mathcal{P}_1(T^2, g)$ jumps from \mathbb{R}^2 to $\{0\}$ as soon as the Gaussian curvature becomes non-identically-zero, and this jump is the sole mechanism by which the bound changes.

Finally, we note that under the pure bidegree hypothesis, exactly one of the two Künneth components is active, so $r - \dim \mathcal{K} \leq 1$. We conjecture (Conjecture 7.2) that the theorem extends to torsion forms of mixed bidegree; this is discussed further in Section 7.

6. ILLUSTRATIVE EXAMPLES

The case $\Sigma_g \times T^2$ (Example of a non-trivial lower bound). Fix $g \geq 2$ and endow Σ_g with any hyperbolic metric and T^2 with the flat metric. The relevant harmonic spaces are

$$\mathcal{H}^1(\Sigma_g) \cong \mathbb{R}^{2g}, \quad \mathcal{H}^2(\Sigma_g) = \text{span}\{\text{vol}_\Sigma\},$$

$$\mathcal{H}^1(T^2) = \text{span}\{dx, dy\}, \quad \mathcal{H}^2(T^2) = \text{span}\{\text{vol}_{T^2}\}.$$

Choose a non-zero harmonic 1-form $\alpha \in \mathcal{H}^1(\Sigma_g)$ (possible because $b_1(\Sigma_g) = 2g > 0$). For the hyperbolic metric on Σ_g ,

$$\mathcal{P}_1(\Sigma_g) = \{0\}, \quad \mathcal{P}_2(\Sigma_g) = \text{span}\{\text{vol}_\Sigma\},$$

since the volume form is parallel, while no harmonic 1-form on a hyperbolic surface is parallel (the holonomy of a hyperbolic metric acts irreducibly on each fibre of $T^*\Sigma_g$, so the only parallel sections of $\Lambda^k T^*\Sigma_g$ are proportional

to the volume form for $k = 2$ and to the zero section for $k = 1$). For the flat metric on T^2 all harmonic forms are parallel:

$$\mathcal{P}_1(T^2) = \mathcal{H}^1(T^2) \cong \mathbb{R}^2, \quad \mathcal{P}_2(T^2) = \mathcal{H}^2(T^2) \cong \mathbb{R}.$$

Consider the cohomology class of pure bidegree $(1, 2)$:

$$[\omega] = \alpha \otimes \text{vol}_{T^2} \in \mathcal{H}^1(\Sigma_g) \otimes \mathcal{H}^2(T^2) \subset H^3(\Sigma_g \times T^2).$$

Since $[\omega]$ lies in a single Künneth summand, the pure bidegree hypothesis of Theorem 5.2 is satisfied for any metric connection whose torsion T represents $[\omega]$ and lies in $\Gamma(V_1^* \otimes \Lambda^2 V_2^*)$. The minimal tensor ranks are

$$r_{2,1} = 0, \quad r_{1,2} = 1, \quad \text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}}) = 1.$$

The obstruction kernel for the unique active component is

$$\ker \tilde{\Psi} = \mathcal{V}_{1,2} \cap (\mathcal{P}_1(\Sigma_g) \otimes \mathcal{P}_2(T^2)) = \{0\},$$

because $\mathcal{P}_1(\Sigma_g) = \{0\}$: the 1-form α is harmonic but not parallel, so the simple tensor $\alpha \otimes \text{vol}_{T^2}$ does not belong to the parallel stratum. Hence $\dim \mathcal{K} = 0$ and the corrected rank is

$$r^{\sharp} = \text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}}) - \dim \mathcal{K} = 1 - 0 = 1.$$

Theorem 5.2 therefore yields the non-trivial lower bound

$$\dim \mathfrak{hol}_p^{\text{off}}(\nabla^C) \geq 1.$$

In particular, any cohomologically calibrated metric connection with torsion of pure bidegree $(1, 2)$ representing $[\omega]$ necessarily has irreducible holonomy across the product splitting $\Sigma_g \times T^2$.

The case $T^2 \times T^2$ (Example of a trivial lower bound). Let $M = T^2 \times T^2$ endowed with the flat product metric $g = g_1 \oplus g_2$. Fix harmonic bases

$$H^1(T^2) = \text{span}\{dx, dy\}, \quad H^2(T^2) = \text{span}\{\text{vol}\}.$$

For the flat metric we have

$$\mathcal{P}_1(T^2) = H^1(T^2) \cong \mathbb{R}^2, \quad \mathcal{P}_2(T^2) = H^2(T^2) \cong \mathbb{R}.$$

The mixed Künneth components are

$$H^2(T^2) \otimes H^1(T^2) \cong \mathbb{R}^2, \quad H^1(T^2) \otimes H^2(T^2) \cong \mathbb{R}^2.$$

Consider the cohomology class

$$[\omega] = \underbrace{\text{vol} \otimes dx}_{(2,1)} + \underbrace{dy \otimes \text{vol}}_{(1,2)}.$$

Its minimal tensor-rank decomposition is

$$\pi_{2,1}([\omega]) = \text{vol} \otimes dx \quad (\text{rank } 1), \quad \pi_{1,2}([\omega]) = dy \otimes \text{vol} \quad (\text{rank } 1).$$

Hence

$$\text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}}) = 1 + 1 = 2.$$

Since $\text{vol} \in \mathcal{P}_2$ and $dx, dy \in \mathcal{P}_1$, both simple tensors lie in the obstruction spaces

$$\ker \Psi_p = \text{span}\{\text{vol} \otimes dx\}, \quad \ker \tilde{\Psi}_p = \text{span}\{dy \otimes \text{vol}\},$$

so $\dim \mathcal{K} = 2$ and

$$r^\sharp = \text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}}) - \dim \mathcal{K} = 2 - 2 = 0.$$

We note that the class $[\omega]$ has non-pure bidegree, since both its $(2, 1)$ and $(1, 2)$ Künneth components are non-trivial, and therefore falls outside the strict hypotheses of Theorem 5.2. However, when $r^\sharp = 0$ the conclusion $\dim \mathfrak{hol}_p^{\text{off}}(\nabla^C) \geq 0$ is universally true and requires no hypothesis on the bidegree of T .

This example illustrates the dichotomy of Remark 5.10: when both factors are flat tori, the parallel-form strata are maximal ($\mathcal{P}_k(T^2) = H^k(T^2)$ for all k) and the obstruction kernel absorbs the entire mixed class, yielding a trivially saturated bound; the geometric interpretation of this regime is discussed in Remark 5.6.

7. CONCLUSION

We have shown that for *admissible* product manifolds $M = M_1 \times M_2$, specifically, products of compact oriented surfaces, and for any cohomologically calibrated metric connection ∇^C with torsion of pure bidegree, the off-diagonal holonomy dimension admits a quantified lower bound on a non-empty open subset $\mathcal{V} \subset M$

$$\dim(\mathfrak{hol}_p^{\text{off}}(\nabla^C)) \geq r^\sharp := \text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}}) - \dim \mathcal{K},$$

where $r = \text{rank}_{\mathbb{R}}([\omega]_{\text{mixed}})$ and $\mathcal{K} = \ker \Psi_p \oplus \ker \tilde{\Psi}_p$ is the sum of the kernels of the linear obstruction maps. The bound is metric-invariant under deformations preserving the parallel-form strata $\mathcal{P}_k(M_i)$. This provides the first *computable* obstruction to reducible holonomy on product manifolds within the admissible class. As shown in Remark 5.10, the behaviour of r^\sharp across the admissible class is completely determined: $r^\sharp = 1$ is forced by topology alone whenever at least one factor has genus $g \geq 2$ and the mixed class is non-trivial, while the transition $r^\sharp = 0 \leftrightarrow 1$ occurs only when a torus factor changes from flat to non-flat.

The examples in Section 6 illustrate both non-trivial and trivial cases, while Remark 2.4 shows by explicit construction that the bound fails to be well-defined without the admissibility hypothesis, establishing sharpness of our assumptions.

When the parallel-form strata absorb the entire mixed class, $r^\sharp = 0$: no irreducible de Rham factor prevents the absorption, and transverse irreducibility (Remark 5.6) is absent. However, $r^\sharp = 0$ does not imply that the off-diagonal holonomy vanishes: as observed in Remark 5.4, the quadratic torsion term $\frac{1}{4}Q_T$ can generate non-trivial off-diagonal curvature even when the harmonic projection ω_h is entirely parallel. This reflects a deliberate

feature of the construction: r^\sharp is defined through the harmonic projection ω_h alone, because for a fixed cohomology class $[\omega]$ infinitely many torsion forms represent $[\omega]$, so any quantity depending on the full torsion would fail to be an invariant of the class. This invariance ensures that r^\sharp is computable from the topology of M and the parallel-form strata of the factor metrics, without knowledge of the specific torsion form.

Strict inequality $\dim \mathfrak{hol}_p^{\text{off}}(\nabla^C) > r^\sharp$ remains open question.

Corollary 7.1 (Extended admissible setting). *Theorem 5.2 extends to products $M_1 \times M_2$ where only the factor contributing the 2-form to the torsion bidegree is required to be a compact oriented surface, while the other factor may be a compact oriented Riemannian manifold of arbitrary dimension. Specifically:*

- (i) *For torsion of pure bidegree $(2, 1)$, M_1 is required to be a compact oriented surface; M_2 may have arbitrary dimension.*
- (ii) *For torsion of pure bidegree $(1, 2)$, M_2 is required to be a compact oriented surface; M_1 may have arbitrary dimension.*

When both factors are surfaces, both bidegree types are covered simultaneously, recovering the full admissible setting of Definition 2.2.

Proof. The proof of Theorem 5.2 for pure bidegree $(2, 1)$ uses the hypothesis $\dim M_1 = 2$ but not $\dim M_2 = 2$; symmetrically, for pure bidegree $(1, 2)$ the proof uses $\dim M_2 = 2$ but not $\dim M_1 = 2$. The only required modification is in *Claim 5.2.1*: the dimension-2 characterisation of symmetric trace-free tensors is replaced by the observation that, when $\nabla^{g_2} \beta_1 \neq 0$, the affine hyperplane $\{Z \in V_2 : \beta_1(Z) = 1\}$ cannot be contained in the kernel of $\nabla^{g_2} \beta_1$ (which is a proper subspace of V_2), so there exists a normalising vector Z for which the evaluation is non-zero; the remainder of the proof proceeds unchanged. \square

In this extended setting, the classification of $\mathcal{P}_1(M_j)$ in Remark 5.10 generalises from genus and Gaussian curvature to the de Rham decomposition theory of the higher-dimensional factor, and the geometric interpretation of r^\sharp (Remark 5.6) acquires additional depth: r^\sharp detects whether the cohomological coupling penetrates beyond the flat sub-factors identified by parallel 1-forms on M_j . A concrete instance arises when $M_2 = N \times S^1$ (respectively $M_1 = N \times S^1$ for bidegree $(1, 2)$): the invariant r^\sharp then predicts whether the irreducibility of the connection survives dimensional reduction along the circle factor. When $r^\sharp = 1$, the torsion 3-form survives restriction to $M_1 \times N$, and the restricted connection remains irreducible on the reduced product; when $r^\sharp = 0$ with $\beta \in \mathcal{P}_1(S^1)$, the torsion vanishes on $M_1 \times N$ and the reduced connection is torsion-free, hence has reducible holonomy on the product.

Conjecture 7.2. *Theorem 5.2 extends to torsion forms of mixed bidegree, i.e. the pure bidegree hypothesis can be removed.*

Under mixed bidegree both Künneth components may be simultaneously active, so $r - \dim \mathcal{K}$ can reach the value 2 in the admissible setting. Establishing the conjecture would require new techniques for the overlap region $\mathcal{W}^{2,1} \cap \mathcal{W}^{1,2}$, where both components of the torsion are simultaneously non-vanishing. A systematic investigation of this extension remains an interesting open problem.

A further direction is to explore extensions beyond the admissible setting. Imposing *Kruskal-rank constraints* on harmonic representatives could yield intrinsic mixed factor spaces even when $b_2(M_i) > 1$, potentially treating Calabi–Yau threefolds or $G_2 \times S^1$ backgrounds. Such extensions would require new algebraic tools and are left for future investigation.

Another further interesting direction for future work could be to explore the implications of these results in the context of deformations of complex structures, a field pioneered by the works of Kodaira [7] and Kodaira–Spencer [8].

Compactness and Hodge theory. Compactness ensures the existence and uniqueness of harmonic representatives [10], and the vanishing of boundary terms in integrations by parts. Extending the results to non-compact manifolds would require an L^2 Hodge framework or alternative analytic hypotheses.

REFERENCES

- [1] N. Aronszajn, *A unique continuation theorem for solutions of elliptic partial differential equations or inequalities of second order*, J. de Math., **T.36** (1957), 235–249.
- [2] Aronszajn, N., Krzywicki, A., Szarski, J. *A unique continuation theorem for exterior differential forms on Riemannian manifolds*. Ark. Mat. 4, 417–453 (1962).
- [3] W. Ambrose and I. M. Singer, *A theorem on holonomy*, Trans. Amer. Math. Soc. **75** (1953), 428–443.
- [4] A. L. Besse, *Einstein Manifolds*, Springer–Verlag, 1987, ISBN: 978-3-540-74120-6.
- [5] R. E. Greene, H. Jacobowitz, *Analytic isometric embeddings*. Ann. of Math. (93) 1 (1971), no. 2, 189–204.
- [6] P. Griffiths and J. Harris, *Principles of Algebraic Geometry*, Wiley–Interscience, New York, 1978, ISBN: 978-0-471-05059-9.
- [7] K. Kodaira, *On Deformations of Complex Analytic Structures, I–II*, Ann. of Math. **67** (1958), 328–466. doi:10.2307/1970055.
- [8] K. Kodaira and D. C. Spencer, *On Deformations of Complex Analytic Structures, III. Stability Theorems for Complex Structures*, Ann. of Math. **70** (1960), 43–76. doi:10.2307/1969838.
- [9] C. B. Morrey Jr., L. Nirenberg, *On the analyticity of the solutions of linear elliptic systems of partial differential equations*. Comm. Pure Appl. Math. 10 (1957), 271–290.
- [10] F. W. Warner, *Foundations of Differentiable Manifolds and Lie Groups*, Springer–Verlag, 1983.