Adiabatic Transport as a Cubical Presheaf: Holonomy, Fibrant Replacement, and Diabatization

Karen Sargsyan
Institute of Chemistry, Academia Sinica
Taipei, Taiwan
karen.sarkisyan@gmail.com

We construct an adiabatic transport presheaf over the cubical singular complex of nuclear configuration space and prove that the projection is a cubical fibration if and only if the Berry curvature vanishes, bridging Cartesian cubical homotopy theory and gauge connection flatness. The cubical face map identities encode holonomy of the Berry connection, ruling out smooth parallel-transport-based fibrant replacement when curvature is nonzero. Applying the algebraic small object argument to a finite computational grid, we show that the resulting fibrant replacement recovers spanning-tree diabatization—the standard algorithm in quantum chemistry—as a canonical, choice-free construction from the model structure. The space of diabatization schemes is the flag manifold $\mathrm{Fl}(\mathbb{C}^N) = \mathrm{U}(N)/\mathrm{U}(1)^N$, whose $\pi_2 \cong \mathbb{Z}^{N-1}$ yields a discrete $(\mathbb{Z}/2)^{N-1}$ invariant classifying conical intersection types on computational grids. We formulate quasi-diabatization as a variational principle with a topological lower bound on residual coupling, and identify the scope of this principle: it is non-trivial for abstract discrete connections and for subspace tracking ($M > N$ active states), but trivially satisfied when the full eigenspace is tracked. Core algebraic properties are independently certified in Cubical Agda.

Keywords: cubical sets, fibrations, Berry connection, holonomy, algebraic small object argument, flag manifold, diabatization

1 Introduction

In quantum chemistry, the Born-Oppenheimer approximation produces electronic states $|\psi_i(R)\rangle$ that depend parametrically on nuclear configuration R. These states form a vector bundle $E \to M$ equipped with the Berry connection [1, 2], a $\mathrm{U}(N)$ gauge field whose holonomy encodes topological features of the electronic structure. At *conical intersections* (CIs)—points where electronic energies become degenerate—the connection has non-trivial holonomy and the adiabatic states become singular.

A *diabatic representation* is a smooth basis that removes these singularities: a gauge-fixing that trivializes the connection. Many methods exist for constructing such bases—Boys localization [3], block diagonalization [4], regularization [5], generalized Mulliken-Hush [6]—and they generally produce different results. This raises a natural categorical question: what is the space of all such gauge-fixings, and how do different methods relate within it?

1.1 Contributions and Categorical Roadmap

We answer this question using the language of cubical fibrations, classifying spaces, and cubical homotopy theory. Our results are organized by six categorical ideas.

Fibrations and flatness. The fibration property (path-independent lifting) corresponds precisely to flatness of the Berry connection (Theorem 2.4). We make this formal by constructing the *adiabatic*

Submitted to:
ACT 2026

transport presheaf $\mathscr{A}(E,\nabla)$ over the cubical singular complex of M, and proving that the projection map is a cubical fibration if and only if the Berry curvature vanishes (Theorem 2.3). The proof reduces to the classical Frobenius integrability theorem; the contribution is the cubical set construction that enables this reduction.

Classifying space. The space of physically distinct diabatic representations is the flag manifold $\mathrm{Fl}(\mathbb{C}^N) = \mathrm{U}(N)/\mathrm{U}(1)^N$ (Theorem 3.2). Each diabatization method corresponds to a point or submanifold, and the topology of this classifying space—specifically $\pi_2 \cong \mathbb{Z}^{N-1}$ (Corollary 3.3)—detects obstructions invisible to single-state Berry phase analysis. (Here π_2 is that of the classifying space $\mathrm{Fl}(\mathbb{C}^N)$, while the Berry phase is holonomy around loops in the base space M; the two invariants live on different spaces.)

Holonomy as cubical coherence. We show (Theorem 5.2) that the abstract cubical face map identities, when instantiated on the transport presheaf, encode the holonomy of the Berry connection. This rules out any smooth parallel-transport-based fibrant replacement when $F \neq 0$ (Corollary 5.5), resolving a natural question about the relationship between Awodey's fibrant replacement and the diabatization algorithm.

ASOA recovers the chemical algorithm. Applying the algebraic small object argument (ASOA) to the transport presheaf on a finite grid, we obtain a canonical fibrant replacement whose transport formula routes through three edges of each plaquette (Proposition 5.6). This is precisely spanning-tree diabatization—the standard algorithm in quantum chemistry (Appendix B)—recovered from pure homotopy theory with no chemical input. The ASOA resolves the spanning-tree ambiguity (exponentially many trees, one canonical resolution) and provides a categorical justification for a procedure that was previously heuristic (Corollary 5.7).

Discrete multi-state invariant. The construction produces a concrete algorithmic output: a discrete topological invariant $\nu \in (\mathbb{Z}/2)^{N-1}$ (Theorem 6.2) that classifies conical intersection configurations from grid data. For $N = 3$, this invariant detects topological structure invisible to individual Berry phases or to the scalar holonomy test $\|W - I\| > \varepsilon$. The computational grid is natively a cubical complex, and the invariant requires no smooth interpolation.

Quasi-diabatization as nearest fibration. When exact diabatization is impossible ($F \neq 0$), we formulate quasi-diabatization as the problem of finding the discrete connection closest to satisfying the cubical fibration condition (Section 8). The resulting variational principle—minimizing the *cubical defect*, derived from the face map identities of Theorem 5.2—has a topological lower bound: the residual coupling cannot vanish on any region enclosing a conical intersection with nonzero ν. We identify the scope of this principle precisely: it is non-trivial for abstract discrete connections and for subspace tracking, but trivially satisfied for eigenvector-overlap connections from a complete eigenspace (Remark 8.2).

Throughout, we distinguish carefully between what is proved and what remains open. The limits of the homotopy-theoretic framework—the flag manifold is a gauge-theoretic, not homotopy-theoretic, invariant—are discussed in Section 9.

2 The Adiabatic Transport Presheaf

We begin with the categorical construction that bridges abstract homotopy theory and gauge theory, as it motivates the geometric results that follow.

2.1 Background: Berry Connection and Holonomy

An N-state electronic system over a nuclear configuration space M determines a rank-N Hermitian vector bundle $\pi : E \to M$ with the Berry connection ∇ [1, 2, 7]. This is a $\mathrm{U}(N)$ gauge field with connection 1-form $A_{ij}(R) = i\langle\psi_i(R)|\nabla_R|\psi_j(R)\rangle$ and curvature $F = dA + A\wedge A$. Parallel transport along a path γ gives a unitary map $P_\gamma : E_{\gamma(0)} \to E_{\gamma(1)}$. The holonomy around a closed loop γ is $W_\gamma = P_\gamma \in \mathrm{U}(E_{\gamma(0)})$.

Proposition 2.1 (Berry phase at conical intersections [1, 8])**.** *For a loop encircling a two-state conical intersection once, $W_\gamma = -I_{2\times 2}$: each state acquires a geometric phase of π.*

2.2 Cubical Singular Complex

For any smooth manifold M, the *cubical singular complex* $S_\square(M)$ is the cubical set with n-cubes

$$S_\square(M)_n = C^\infty(I^n, M), \tag{1}$$

where $I = [0,1]$. Face maps $\partial_k^\varepsilon : S_\square(M)_n \to S_\square(M)_{n-1}$ restrict to the face $\{x_k = \varepsilon\}$ for $\varepsilon \in \{0,1\}$, and degeneracy maps s_k insert a constant direction. This is a standard construction: the cubical analogue of the simplicial singular complex.

2.3 The Construction

Definition 2.2 (Adiabatic transport presheaf)**.** *Let $(E,\nabla) \to M$ be a smooth Hermitian vector bundle with unitary connection. The* adiabatic transport presheaf *$\mathscr{A}(E,\nabla)$ is the cubical set with n-cubes*

$$\mathscr{A}_n = \{(\sigma, s) : \sigma \in C^\infty(I^n, M),\ s \in \Gamma(\sigma^* E),\ (\sigma^*\nabla)_{\partial_k} s = 0 \textit{ for all } k = 1, \ldots, n\}, \tag{2}$$

i.e., pairs of a smooth n-cube in M and a section of the pullback bundle that is covariantly constant in every coordinate direction simultaneously. Face and degeneracy maps are induced by restriction and insertion: if $(\sigma, s) \in \mathscr{A}_n$, then restricting to the face $\{x_k = \varepsilon\}$ gives $(\partial_k^\varepsilon \sigma, s|_{\{x_k=\varepsilon\}}) \in \mathscr{A}_{n-1}$ since a section parallel in n directions restricts to one parallel in the remaining $n-1$ directions; inserting a constant k-th coordinate gives $(s_k\sigma, s \circ \mathrm{pr}) \in \mathscr{A}_{n+1}$ since the added direction contributes $\nabla_{\partial_k} s = 0$ trivially. The cubical identities are inherited from $S_\square(M)$.

The projection $p : \mathscr{A}(E,\nabla) \to S_\square(M)$ forgets the section s. At low dimensions:

- $\mathscr{A}_0 = \{(x, v) : x \in M, v \in E_x\}$—points of the total space E.
- $\mathscr{A}_1 = \{(\gamma, s) : \gamma : I \to M, \nabla_{\dot\gamma} s = 0\}$—parallel sections along paths. Each is determined by (γ, v) where $v = s(0) \in E_{\gamma(0)}$, via parallel transport.
- $\mathscr{A}_2 = \{(\sigma, s) : \sigma : I^2 \to M, \nabla_{\partial_1} s = 0, \nabla_{\partial_2} s = 0\}$—sections covariantly constant in both directions over a square.

2.4 The Bridge Theorem

Theorem 2.3 (Adiabatic transport presheaf and fibrations)**.** *Let (E,∇) be a smooth Hermitian vector bundle with unitary connection and curvature 2-form F over a manifold M. The projection $p : \mathscr{A}(E,\nabla) \to S_\square(M)$ satisfies:*

1. (1-dimensional filling) *For every path $\gamma : I \to M$ and every $v \in E_{\gamma(0)}$, there exists a unique $(\gamma, s) \in \mathscr{A}_1$ with $s(0) = v$.*

2. (2-dimensional filling) *For a smooth map* $\sigma : I^2 \to M$ *and any* $v \in E_{\sigma(0,0)}$, *a section* $s \in \Gamma(\sigma^* E)$ *satisfying* $s(0,0) = v$ *and* $(\sigma^*\nabla)_{\partial_k} s = 0$ *for* $k = 1,2$ *exists for every such* v *if and only if* $\sigma^* F = 0$.
3. (Fibration characterization) *The map* p *satisfies the open box filling condition for all* n*-dimensional boxes if and only if* $F = 0$ *on* M.

Proof. (1): Existence and uniqueness of s follow from the Picard-Lindelöf theorem applied to the ODE $\nabla_{\dot\gamma} s = 0$ with initial condition $s(0) = v$.

(2): The system $\nabla_{\partial_1} s = 0$, $\nabla_{\partial_2} s = 0$ is a system of first-order PDEs for $s : I^2 \to \sigma^* E$. The Frobenius integrability condition requires the commutator to vanish:

$$[\nabla_{\partial_1}, \nabla_{\partial_2}]\, s = F(\sigma_*\partial_1, \sigma_*\partial_2)\, s = (\sigma^* F)\, s = 0. \tag{3}$$

If $\sigma^* F = 0$, the system is integrable and a unique solution exists for any initial value $s(0,0) = v$ (Frobenius theorem). Conversely, if $\sigma^* F \neq 0$ at some point $x_0 \in I^2$, then for any $v \in E_{\sigma(0,0)}$ whose parallel transport to x_0 does not lie in $\ker F(x_0)$—a proper subspace when $F(x_0) \neq 0$—no solution exists on all of I^2. (The zero section is always a solution; the point is that nonzero initial data generically yields no global solution.)

(3): An open box in dimension n provides boundary data on all but one face of I^n. Filling requires extending to a section s on I^n with $\nabla_{\partial_k} s = 0$ for all k, matching this data.

($\Leftarrow$) If $F = 0$, the system $\nabla_{\partial_k} s = 0$ for $k = 1, \ldots, n$ is integrable for every σ, and a unique solution exists given initial data $s(c) = v$ at any corner c. The boundary data consists of parallel sections on each face, sharing corner values. Since $F = 0$ makes parallel transport path-independent, the Frobenius solution from any corner agrees with the boundary data on each face (both are determined by the same corner value via path-independent transport).

($\Rightarrow$) If $F(x_0) \neq 0$ for some $x_0 \in M$, choose $\sigma : I^2 \to M$ with $\sigma^* F \neq 0$. The faces of the open box carry nonzero parallel sections (existing by (1)). Any filler must simultaneously satisfy $\nabla_{\partial_1} s = 0$ and $\nabla_{\partial_2} s = 0$ while matching this boundary data, but by (2), for a generic choice of initial vector $v \notin \ker F(x_0)$, no simultaneously parallel section with the prescribed boundary data exists over this σ. Higher-dimensional boxes reduce to the 2-dimensional case by restricting to 2-faces.

The mathematical content is the Frobenius integrability theorem; the contribution is Definition 2.2, which translates the PDE condition into the cubical fibration property. □

In Awodey's framework [9], p is a *biased fibration* if it has the right lifting property against open box inclusions (Definition 4.1, op. cit.); biased and unbiased fibration conditions coincide in the Cartesian cubical model structure (Theorem 9.8, op. cit.), so our characterization extends to both.[1]

Combining Theorem 2.3 with classical results yields:

Theorem 2.4 (Detection: fibrations, flatness, and diabatic bases)**.** *On a simply connected domain* $U \subseteq M$, *the following are equivalent:* (1) $F = 0$ *on* U; (2) *parallel transport is path-independent on* U (*Ambrose-Singer* [11])*;* (3) *the adiabatic transport presheaf* $p : \mathscr{A} \to S_\square(U)$ *satisfies the open box filling condition* (*Theorem* 2.3(3))*;* (4) *a global smooth diabatic basis exists on* U ([12]).

Condition (3) is the new content: it identifies the obstruction to diabatization with the failure of a cubical presheaf to be a fibration.

---

[1] Awodey shows [9, p. 5] that the purely Cartesian cubical model structure is not Quillen equivalent to spaces; adding an equivariance condition restores the equivalence [10]. Our π_2 computations use classical homotopy theory; the cubical framework provides the formal bridge.

3 The Flag Manifold Classification

We now classify the space of all diabatic representations—the gauge-fixings that trivialize the Berry connection—and compute its homotopy type.

Definition 3.1 (Diabatic representation)**.** *A* diabatic representation *is an ordered orthonormal basis* $\{|\phi_1\rangle, \dots, |\phi_N\rangle\}$ *for* $E_{R_0} \cong \mathbb{C}^N$ *at a reference configuration* R_0*, extended to a neighborhood by parallel transport along a chosen system of paths from* R_0 *(when* $F = 0$*, this extension is path-independent). Two bases are* physically equivalent *if they differ by independent phases:* $|\phi_i\rangle \sim e^{i\theta_i}|\phi_i\rangle$*.*

Theorem 3.2 (Classification)**.** *The space of physically distinct diabatic representations for an N-state system is the complete flag manifold:*

$$\mathrm{Fl}(\mathbb{C}^N) = \mathrm{U}(N)/\mathrm{U}(1)^N. \tag{4}$$

Proof. Any two ordered orthonormal bases for $\mathbb{C}^N$ are related by $U \in \mathrm{U}(N)$. The physical equivalence $|\phi_i\rangle \sim e^{i\theta_i}|\phi_i\rangle$ identifies U with UD for any diagonal unitary $D \in \mathrm{U}(1)^N$. The quotient $\mathrm{U}(N)/\mathrm{U}(1)^N$ is the complete flag manifold: a coset $U \cdot \mathrm{U}(1)^N$ determines the flag $L_k = \mathrm{span}\{Ue_1, \dots, Ue_k\}$, and conversely a complete flag $0 \subset L_1 \subset \cdots \subset L_N = \mathbb{C}^N$ determines a basis up to phases. □

Corollary 3.3 (Topological invariants)**.** *The flag manifold satisfies* $\pi_1(\mathrm{Fl}(\mathbb{C}^N)) = 0$ *and* $\pi_2(\mathrm{Fl}(\mathbb{C}^N)) \cong \mathbb{Z}^{N-1}$*.*

Proof. From the fibration $\mathrm{U}(1)^N \to \mathrm{U}(N) \to \mathrm{Fl}(\mathbb{C}^N)$, the long exact sequence gives

$$\pi_2(\mathrm{U}(N)) \to \pi_2(\mathrm{Fl}(\mathbb{C}^N)) \to \pi_1(\mathrm{U}(1)^N) \to \pi_1(\mathrm{U}(N)).$$

Since $\pi_2(\mathrm{U}(N)) = 0$, we obtain $\pi_2(\mathrm{Fl}(\mathbb{C}^N)) \cong \ker(\mathbb{Z}^N \xrightarrow{\Sigma} \mathbb{Z}) \cong \mathbb{Z}^{N-1}$, where Σ sends $(n_1, \dots, n_N) \mapsto \sum n_i$. Surjectivity of Σ yields $\pi_1(\mathrm{Fl}(\mathbb{C}^N)) = 0$. □

The significance of these invariants is that $\pi_2 = \mathbb{Z}^{N-1}$ detects topological structure invisible to single-state Berry phase analysis. For $N = 2$, $\mathrm{Fl}(\mathbb{C}^2) = \mathbb{CP}^1 \cong S^2$ and $\pi_2 = \mathbb{Z}$ recovers the single Berry phase invariant. For $N \geq 3$, the additional $\mathbb{Z}$-summands detect independent obstructions arising from distinct types of conical intersection.

N	$\mathrm{Fl}(\mathbb{C}^N)$	$\dim_{\mathbb{R}}$	π_2	χ
2	$S^2 \cong \mathbb{CP}^1$	2	$\mathbb{Z}$	2
3	$\mathrm{Fl}(\mathbb{C}^3)$	6	$\mathbb{Z}^2$	6
4	$\mathrm{Fl}(\mathbb{C}^4)$	12	$\mathbb{Z}^3$	24
N	$\mathrm{Fl}(\mathbb{C}^N)$	$N(N-1)$	$\mathbb{Z}^{N-1}$	$N!$

Each diabatization method (parallel transport, Boys localization, block diagonalization, etc.) corresponds to a point or submanifold of $\mathrm{Fl}(\mathbb{C}^N)$; the assignment respects bundle isomorphisms: a connection-preserving isomorphism $(E, \nabla) \xrightarrow{\sim} (E', \nabla')$ induces, via the fiber map at the reference point, a diffeomorphism $\mathrm{Fl}(\mathbb{C}^N) \to \mathrm{Fl}(\mathbb{C}^N)$ mapping each method's locus to itself. The explicit mapping for the principal methods in quantum chemistry, with a summary table, is given in Appendix A.

4 Three-State Systems and π_2 Invariants

The flag manifold classification reveals new topological structure starting at $N = 3$, where $\pi_2 = \mathbb{Z}^2$ provides two independent invariants rather than the single Berry phase available for two-state systems.

4.1 Physical Interpretation of $(\mathbb{Z}/2)^2$

Consider a three-state system with two types of conical intersection: $\mathrm{CI}_{1\text{-}2}$ (states 1 and 2 degenerate) and $\mathrm{CI}_{2\text{-}3}$ (states 2 and 3 degenerate). The holonomy eigenphases $\gamma_1, \gamma_2, \gamma_3$ (assigned to adiabatic states by continuous tracking from a basepoint) satisfy $\gamma_1 + \gamma_2 + \gamma_3 \equiv 0 \pmod{2\pi}$ (since each isolated CI contributes $-I$ to a 2×2 block, giving $\det W_\gamma = \prod_k e^{i\gamma_k} = 1$). For isolated conical intersections, each γ_k is 0 or π. Defining the *Berry parity*

$$b_k = \gamma_k/\pi \bmod 2 \in \mathbb{Z}/2, \tag{5}$$

the constraint becomes $b_1 + b_2 + b_3 = 0$ in $\mathbb{Z}/2$, leaving two independent invariants. A convenient choice of generators for $(\mathbb{Z}/2)^2$ is

$$\nu_1 = b_1, \qquad \nu_2 = b_1 + b_2 \bmod 2. \tag{6}$$

Each $\mathbb{Z}/2$-summand detects a different conical intersection type.

To relate this to $\pi_2(\mathrm{Fl}(\mathbb{C}^N)) \cong \mathbb{Z}^{N-1}$: the eigenphase invariant captures the mod-2 image of the full $\mathbb{Z}^{N-1}$ Chern number classification, since both Chern numbers $c_k = +1$ and $c_k = -1$ produce the same holonomy eigenvalue $e^{i\pi} = -1$. For generic isolated CIs (topological charge ±1), this mod-2 detection is complete.

4.2 Numerical Verification

We verify the $(\mathbb{Z}/2)^2$ structure using a three-state model Hamiltonian with two CIs at distinct locations (Appendix C). Computing holonomy eigenphases around three topologically distinct loops—enclosing $\mathrm{CI}_{1\text{-}2}$ alone (loop A), $\mathrm{CI}_{2\text{-}3}$ alone (loop B), or both (loop C)—the group law $\nu(C) = \nu(A) + \nu(B)$ holds in $(\mathbb{Z}/2)^2$, with each CI type mapping to a different generator. State 2 acquires Berry parity $b_2 = 1$ from each CI independently, so encircling both yields $b_2 = 1 + 1 = 0 \bmod 2$: the Berry phase of state 2 alone cannot distinguish the two-CI configuration from the trivial one, but the $(\mathbb{Z}/2)^2$ invariant $(\nu_1, \nu_2) = (1,1) \neq (0,0)$ does.

While the homotopy groups $\pi_2(\mathrm{Fl}(\mathbb{C}^N))$ are standard in algebraic topology, their interpretation as invariants classifying multi-state conical intersection patterns—and the demonstration that $(\mathbb{Z}/2)^{N-1}$ resolves ambiguities invisible to individual Berry phases—appears to be new.

5 Resolution of the Fibrant Replacement Problem

Having established the classifying space and its π_2 invariants, we return to the fibrant replacement question raised by Theorem 2.3: when $F \neq 0$, what does the ASOA produce?

We originally conjectured (in an earlier version of this paper) that Awodey's algebraic small object argument, applied to $p : \mathscr{A} \to S_\square(M)$, would recover the parallel transport trivialization. We now show this conjecture is false when $F \neq 0$, and that the obstruction is precisely the holonomy of the connection—encoded, remarkably, in the cubical face map identities.

5.1 The Natural Candidate

The parallel transport trivialization suggests the following candidate for a fibrant replacement. Fix a smooth n-cube $\sigma : I^n \to M$. Since every 1-cube in $\mathscr{A}$ is determined by its initial value (via the ODE for parallel transport), the simplest enlargement of $\mathscr{A}$ that "fills all boxes" would be to drop the simultaneous-integrability condition and track only the initial value:

Definition 5.1 (Transport candidate)**.** *Let* $\mathbf{0} = (0,\dots,0) \in I^n$. *Define* $\mathscr{T}_n = \{(\sigma, v) : \sigma \in C^\infty(I^n, M),\, v \in E_{\sigma(\mathbf{0})}\}$ *with face maps* $\partial_k^0(\sigma, v) = (\partial_k^0\sigma,\, v)$, $\partial_k^1(\sigma, v) = (\partial_k^1\sigma,\, \tau_{\sigma;k}(v))$, *where* $\tau_{\sigma;k}$ *is parallel transport from* $E_{\sigma(\mathbf{0})}$ *to* $E_{\sigma(\mathbf{0}+\mathbf{e}_k)}$ *along* $t \mapsto \sigma(0,\dots,t,\dots,0)$.

This definition is modeled on the diabatization algorithm: each n-cube carries an initial frame v at the corner $\mathbf{0}$, propagated outward by parallel transport along coordinate directions.

5.2 The Holonomy Obstruction

Theorem 5.2 (Cubical identities encode holonomy)**.** *The transport candidate* $\mathscr{T}$ *of Definition* 5.1 *is a cubical set if and only if* $F = 0$.

Proof. The cubical identity $\partial_j^\varepsilon \partial_k^\delta = \partial_{k-1}^\delta \partial_j^\varepsilon$ (for $j < k$) must hold for all face maps. Consider $n = 2$, $j = 1$, $k = 2$, $\varepsilon = \delta = 1$, applied to (σ, v) with $\sigma : I^2 \to M$ and $v \in E_{\sigma(0,0)}$.

Left-hand side: $\partial_1^1\,\partial_2^1\,(\sigma, v)$. First, ∂_2^1 transports v along the left edge $t \mapsto \sigma(0,t)$, giving $\tau_{\text{left}}(v) \in E_{\sigma(0,1)}$. Then ∂_1^1 transports the result along the top edge $t \mapsto \sigma(t,1)$, giving $\tau_{\text{top}} \circ \tau_{\text{left}}(v) \in E_{\sigma(1,1)}$.

Right-hand side: $\partial_1^1\,\partial_1^1\,(\sigma, v)$ (note: $j = 1, k = 2$ gives $\partial_{k-1}^\delta \partial_j^\varepsilon = \partial_1^1\partial_1^1$; the outer ∂_1^1 acts on the 2-cube, the inner on the resulting 1-cube). First, ∂_1^1 on the 2-cube transports v along the bottom edge $t \mapsto \sigma(t,0)$, giving $\tau_{\text{bottom}}(v) \in E_{\sigma(1,0)}$. Then ∂_1^1 (now acting on the resulting 1-cube) transports along the right edge $t \mapsto \sigma(1,t)$, giving $\tau_{\text{right}} \circ \tau_{\text{bottom}}(v) \in E_{\sigma(1,1)}$.

The identity $\partial_1^1\,\partial_2^1 = \partial_1^1\,\partial_1^1$ requires

$$\tau_{\text{top}} \circ \tau_{\text{left}} = \tau_{\text{right}} \circ \tau_{\text{bottom}} \tag{7}$$

for all smooth maps $\sigma : I^2 \to M$ and all $v \in E_{\sigma(0,0)}$.

This is the statement that holonomy around every contractible loop in M is trivial: equivalently, $\tau_{\text{left}}^{-1} \circ \tau_{\text{top}}^{-1} \circ \tau_{\text{right}} \circ \tau_{\text{bottom}}$ around $\partial(\sigma(I^2))$ (starting from $E_{\sigma(0,0)}$) equals the identity. By the Ambrose-Singer theorem, this holds for all σ if and only if $F = 0$.

For the converse: when $F = 0$, all parallel transports commute in the required sense, all cubical identities hold, and $\mathscr{T} = \mathscr{A}$ (every initial value extends uniquely to a simultaneously parallel section). The remaining face map identities (those involving at least one $\varepsilon = 0$ face) hold unconditionally, since ∂_k^0 does not involve parallel transport. □

Remark 5.3. *The result does not depend on the choice of corner* $\mathbf{0}$ *or the convention for coordinate-axis paths. Any candidate of the form "track a fiber element at a corner, propagate by parallel transport along chosen paths to define face maps" encounters the same obstruction: the cubical identity* (7) *for any pair of paths sharing endpoints encodes holonomy around the enclosed region.*

Corollary 5.4 (Discrete holonomy obstruction)**.** *Let* $\mathscr{G}$ *be a finite grid with* $\mathrm{U}(N)$*-valued overlap matrices* $\{U_e\}$, *and define the discrete transport candidate by assigning to each vertex* k *a fiber element* $v_k \in E_k$, *with face maps* $\partial_e^0(k, v_k) = (k, v_k)$ *and* $\partial_e^1(k, v_k) = (k', U_e \cdot v_k)$ *for each edge* $e = (k, k')$. *This candidate satisfies the cubical identities if and only if* $W_f = I$ *for every plaquette* f.

Proof. The identity $\partial_1^1\partial_2^1 = \partial_1^1\partial_1^1$ at plaquette f requires $\tau_{\text{top}} \circ \tau_{\text{left}} = \tau_{\text{right}} \circ \tau_{\text{bottom}}$, i.e., $W_f = I$. Conversely, $W_f = I$ for all f implies all identities involving $\varepsilon = \delta = 1$ faces hold; those with at least one $\varepsilon = 0$ face hold unconditionally. □

This is the self-contained discrete analogue of Theorem 5.2: no smooth theory or Ambrose-Singer is needed; each plaquette is checked directly. The Cubical Agda formalization (Appendix D) certifies exactly this equivalence for $\mathbb{Z}/2$-valued connections.

5.3 Consequences for Fibrant Replacement

Corollary 5.5. *When $F \neq 0$, the algebraic small object argument fibrant replacement $\hat{\mathscr{A}}$ of $p : \mathscr{A} \to S_\square(M)$ is not naturally isomorphic to any cubical set whose face maps are defined by parallel transport along smooth paths. The ASOA necessarily introduces formal (non-geometric) fillers at dimension ≥ 2.*

Proof. The ASOA fibrant replacement $\hat{\mathscr{A}}$ is a cubical set by construction (built by pushouts in cSet). By Theorem 5.2, any cubical set whose face maps use parallel transport satisfies the cubical identities only if $F = 0$, in which case $\mathscr{A}$ is already fibrant and the ASOA is trivial. □

5.4 The Positive Content

While the conjecture is false, the analysis yields positive results. First, Theorem 5.2 reveals that the cubical face map identities encode the holonomy conditions for the Berry connection: the cubical combinatorics does not merely translate the differential geometry—it *is* the differential geometry. Second, when $F = 0$ on a simply connected domain, $\mathscr{T} = \mathscr{A}$ and p is already a fibration. The ASOA returns $\mathscr{A}$ itself: the parallel transport trivialization is the canonical fibrant replacement, answering the canonicity question from the chemistry literature in the affirmative for the flat case. Third, the non-flat ASOA on a finite grid recovers the standard chemical algorithm (Corollary 5.7 below).

5.5 The ASOA on a Finite Grid

We now apply the ASOA to the *restricted* map $p|_{\mathscr{G}} : \mathscr{A}|_{\mathscr{G}} \to \mathscr{G}$, where $\mathscr{G}$ is a finite cubical complex with $\mathrm{U}(N)$ overlap matrices $\{U_e\}$ and plaquette holonomies $W_f = \prod_{e \in \partial f} U_e$. This is distinct from restricting the abstract ASOA fibrant replacement of $p : \mathscr{A} \to S_\square(M)$ to $\mathscr{G}$; however, for a 2-dimensional grid, $\mathscr{G}$ has no cells above dimension 2 and the ASOA terminates after that stage, making the finite construction self-contained. The ASOA operates by dimension. At dimension 1, all open boxes are filled (edge transport exists). At dimension 2, consider a plaquette f with edge transports $\tau_B, \tau_L, \tau_R, \tau_T$. For the open box missing the top face with initial $v \in E_{v_{00}}$, the three given faces are bottom, left, and right transport from v. A filler needs a top 1-cube from $\tau_L(v)$ to $\tau_R \tau_B(v)$, but geometric transport gives $\tau_T \tau_L(v) \neq \tau_R \tau_B(v)$ when $W_f \neq I$.

Proposition 5.6 (ASOA transport formula)**.** *The ASOA resolves the open box by attaching a formal 2-cube whose top face implements transport $\hat{\tau}_T : E_{v_{01}} \to E_{v_{11}}$:*

$$\hat{\tau}_T = \tau_R \circ \tau_B \circ \tau_L^{-1} = \tau_T \circ \mathrm{Ad}_{\tau_L}(W_f), \tag{8}$$

where $\mathrm{Ad}_{\tau_L}(W_f) = \tau_L W_f \tau_L^{-1}$. The ASOA transport differs from geometric transport by the holonomy conjugated to the source fiber.

Proof. The first equality is the definition: the filler maps $\tau_L(v) \mapsto \tau_R \tau_B(v)$. For the second, use $W_f = \tau_L^{-1} \tau_T^{-1} \tau_R \tau_B$ to get $\tau_R \tau_B = \tau_T \tau_L W_f$, hence $\hat{\tau}_T = \tau_T \tau_L W_f \tau_L^{-1}$. □

The formula (8) has a striking consequence: the ASOA transport $\hat{\tau}_T = \tau_R \circ \tau_B \circ \tau_L^{-1}$ routes through the three non-missing edges, which is exactly what a spanning-tree diabatization does when the missing edge is the cut edge.

Corollary 5.7 (ASOA recovers spanning-tree diabatization)**.** *On a simply connected 2D grid, the ASOA fibrant replacement* $\hat{\mathscr{A}}|_{\mathscr{G}}$ *produces, for each open box type* (*missing top, bottom, left, or right*)*, a flat connection whose gauge equivalence class coincides with the spanning-tree diabatization that cuts the corresponding edge family. In particular, the ASOA canonically recovers the standard chemical algorithm* (*Appendix* B) *from the model structure, with no chemical input.*

Proof. Proposition 5.6 gives $\hat{\tau}_T = \tau_R \tau_B \tau_L^{-1}$, which is the parallel transport along the spanning tree that includes bottom, left, and right edges and cuts the top. By Theorem E.1 (Appendix E), the gauge class is determined by $\{W_f\}$, and both constructions produce the same gauge class. □

Spanning-tree diabatization, used in quantum chemistry since the 1990s [4] and motivated by physical reasoning, is recovered by the ASOA from the model structure on cubical sets [9] with no chemical input. The ASOA resolves the spanning-tree ambiguity: for a grid with $|F|$ plaquettes, there are exponentially many spanning trees but exactly one ASOA resolution per open box type.

On a simply connected 2D grid, the gauge equivalence class of $\hat{\mathscr{A}}|_{\mathscr{G}}$ is determined by the plaquette holonomies $\{W_f\}$ alone (Appendix E), just as for any spanning-tree diabatization. However, the specific ASOA representative carries more: the correction $\mathrm{Ad}_{\tau_L}(W_f)$ at each plaquette encodes how the holonomy defect appears in the local gauge frame, and this gauge-covariant structure depends on the full connection, not just $\{W_f\}$ (Appendix E, Eq. (20)).

The homotopy type of $\hat{\mathscr{A}}|_{\mathscr{G}}$ is trivial (the fibers $E_x \cong \mathbb{C}^N$ are contractible), confirming the sharp limit of Section 9. The value lies in the algebraic structure: the number, location, and magnitude of formal fillers encode the spatial pattern of Berry curvature on the grid.

Remark 5.8 (AWFS structure and initiality)**.** *The ASOA factorization* $\mathscr{A}|_{\mathscr{G}} \xrightarrow{j} \hat{\mathscr{A}}|_{\mathscr{G}} \xrightarrow{q} \mathscr{G}$ *is an instance of the algebraic weak factorization system (AWFS) of Garner [13]. The fibration side* q *carries* algebra structure*: a functorial choice of filler for every open box. On our grid, this algebra structure is precisely the ASOA transport formula* (8)*—a chosen filler* $\hat{\tau}_T = \tau_R \tau_B \tau_L^{-1}$ *at each plaquette per box type.*

The key categorical property is initiality*:* $\hat{\mathscr{A}}|_{\mathscr{G}}$ *is initial among fibrant replacements of* $\mathscr{A}|_{\mathscr{G}}$ *equipped with chosen fillers. Any other such replacement receives a unique filler-preserving map from* $\hat{\mathscr{A}}|_{\mathscr{G}}$ *(see [13, Thm. 3.1]; the general framework is surveyed in [14, Ch. 12]). This is the precise sense of "canonical": not merely a particular choice, but the universal one—canonicity via a universal property, as the fundamental groupoid is the universal monodromy target. The exponential spanning-tree ambiguity is resolved because spanning-tree diabatizations produce fillers that are geometric but non-canonical (tree-dependent), while the ASOA produces fillers that are equally concrete (Proposition 5.6) but canonically determined by the universal property.*

The ASOA is also compatible with grid refinement. When $\mathscr{G}$ *is subdivided into* $\mathscr{G}'$*, each plaquette* f *of* $\mathscr{G}$ *is partitioned into sub-plaquettes, and the holonomy factorizes:* $W_f = \prod_{f' \subset f} W_{f'}$ *(up to conjugation). Since the ASOA on a simply connected 2D grid is determined up to gauge by plaquette holonomies (Theorem E.1), the gauge class of* $\hat{\mathscr{A}}|_{\mathscr{G}'}$ *is compatible with that of* $\hat{\mathscr{A}}|_{\mathscr{G}}$*: the coarse gauge class is determined by the coarse plaquette holonomies, which are products of the fine ones. This compatibility underlies the* $O(h^2)$ *convergence noted in Section 7.2.*

6 A Discrete Multi-State Topological Invariant

The $\pi_2 = \mathbb{Z}^{N-1}$ classification (Section 4) provides topological invariants for multi-state CI configurations. We now show that the discrete holonomy obstruction (Corollary 5.4) leads to an algorithm for computing these invariants directly from grid data, extracting information the standard algorithm currently discards.

The standard topology check (Appendix B, Step 2) computes the plaquette holonomy $W_f = \prod_{e \in \partial f} U_e$ and flags f as enclosing a CI when $\|W_f - I\| > \varepsilon$—a binary test. But the eigenvalue structure of W_f, tracked by state label, encodes the $(\mathbb{Z}/2)^{N-1}$ type.

Definition 6.1 (Multi-state Berry numbers on a grid)**.** *Let γ be a closed loop in a grid $\mathscr{G}$ with* $\mathrm{U}(N)$*-valued overlap matrices, bounding a simply connected region Σ. The holonomy $W_\gamma = \prod_{e \in \gamma} U_e$ has eigenvalues $e^{i\gamma_1}, \dots, e^{i\gamma_N}$, assigned to adiabatic states by tracking eigenvectors from a basepoint on γ. The* multi-state Berry numbers *are*

$$\nu(\Sigma) = (\nu_1, \dots, \nu_{N-1}) \in (\mathbb{Z}/2)^{N-1}, \qquad \nu_1 = b_1, \quad \nu_j = b_1 + \cdots + b_j \bmod 2, \tag{9}$$

where $b_k = \gamma_k/\pi \bmod 2 \in \mathbb{Z}/2$.

Theorem 6.2 (Discrete CI classification)**.** *The multi-state Berry numbers satisfy:* (a) *gauge invariance (eigenvalues of W_γ are conjugation-invariant);* (b) *additivity across disjoint regions in $(\mathbb{Z}/2)^{N-1}$;* (c) *the map ν realizes the mod-2 image of $\pi_2(\mathrm{Fl}(\mathbb{C}^N)) \cong \mathbb{Z}^{N-1}$, with generators $\mathbf{e}_j$ corresponding to elementary $(j, j{+}1)$-CIs. For isolated CIs of topological charge ± 1, ν completely classifies the CI type.*

Proof. Parts (a)–(b) follow from properties of the holonomy product (gauge invariance of eigenvalues; multiplicativity of determinants for composed loops).

Part (c): On a loop $\gamma = \partial\Sigma$ avoiding all degeneracies, the eigenvalues of the parameter-dependent Hamiltonian are distinct, so the bundle splits smoothly as $E|_\Sigma \cong L_1 \oplus \cdots \oplus L_N$ away from the CI points in Σ, where each L_k is the k-th eigenspace line bundle (see [11, Ch. 4]). The Berry connection restricts to a $\mathrm{U}(1)$ connection on each L_k, with holonomy eigenvalue $e^{i\gamma_k}$. For an isolated CI of charge ± 1 between states j and $j{+}1$, the local normal form gives monodromy $e^{i\gamma_j} = e^{i\gamma_{j+1}} = -1$ (a π phase) with all other eigenvalues equal to $+1$ [12]. Define the *signed CI count* $c_k \in \mathbb{Z}$ as the algebraic number of CIs in Σ involving state k, with the Chern sign convention: each charge-± 1 CI between states $j, j{+}1$ contributes ± 1 to c_j and ∓ 1 to c_{j+1}. Then $\gamma_k \equiv \pi c_k \pmod{2\pi}$ (since $e^{i\pi} = e^{-i\pi} = -1$), so $e^{i\gamma_k} = e^{i\pi c_k}$ and $b_k = c_k \bmod 2$. Since each CI contributes opposite signs to the two states involved, $\sum_k c_k = 0$, giving $\mathbf{c} = (c_1, \dots, c_N) \in \ker(\mathbb{Z}^N \to \mathbb{Z}) \cong \mathbb{Z}^{N-1}$, which maps to ν under mod-2 reduction. When each pair of adjacent states has at most one CI of charge ± 1, the entries satisfy $|c_k| \leq 1$; since $e^{i\pi} = e^{-i\pi}$, the holonomy eigenvalues cannot distinguish charge $+1$ from charge -1, and the mod-2 reduction loses no holonomy-distinguishable information. □

Remark 6.3 (Higher-charge conical intersections)**.** *A charge-± 2 CI produces monodromy $e^{2i\pi} = +1$, giving $b_k = 0$: such CIs are invisible to the mod-2 invariant ν. The full $\mathbb{Z}^{N-1}$ Chern classification detects them ($c_k = \pm 2$), but requires smooth interpolation. In practice this is rarely a limitation: charge-± 2 CIs are generically unstable under perturbation and split into pairs of charge-± 1 CIs, which ν detects individually.*

As demonstrated in Section 4, the scalar holonomy test $\|W - I\| > \varepsilon$ cannot distinguish a loop encircling two CIs (with $\nu = (1,1)$) from one encircling a single CI ($(1,0)$ or $(0,1)$); the $(\mathbb{Z}/2)^2$ invariant resolves this (Appendix C).

The discrete setting is essential: computing Chern numbers smoothly requires interpolation, surface choice, and gauge singularity handling. The discrete ν avoids all three—overlap matrices are the primary data, grid plaquettes the canonical 2-cells, and eigenvalue extraction is gauge-invariant by construction. Gauge invariance, additivity, and the fibration–flatness equivalence have been independently verified by a Cubical Agda formalization of the $\mathbb{Z}/2$ Berry parity model (Appendix D). The additional cost beyond the standard algorithm is one $N \times N$ eigenvalue decomposition per loop.

7 The Cubical-Gauge Dictionary

We collect the correspondences established in Sections 2–6 into a single dictionary, and discuss what the cubical formulation adds beyond recasting classical differential geometry.

7.1 The Complete Dictionary

Cubical type theory	**Berry connection**	**Reference**
Cubical set X	Sing. complex $S_\square(M)$	Def.
Map $p : \mathscr{A} \to S_\square(M)$	Transport presheaf	Def. 2.2
Open box filling	Frobenius integrability	Thm. 2.3
p fibration $\Leftrightarrow F = 0$	Flatness detection	Thm. 2.4
Cubical face identities	Holonomy triviality	Thm. 5.2
Discrete face identities	$W_f = I$ for all plaquettes	Cor. 5.4
Fibrant replacement	Not parallel transport	Cor. 5.5
ASOA on finite grid	Spanning-tree diabatization	Cor. 5.7
Discrete holonomy on grid	Multi-state $\nu \in (\mathbb{Z}/2)^{N-1}$	Thm. 6.2
Nearest fibrant object	Quasi-diabatization principle	§8

All rows except the last correspond to results established above; the final row is the subject of Section 8.

7.2 Computational Grids Are Natively Cubical

Diabatization algorithms operate on *grids* of nuclear configurations: a set of points $\{R_{i_1,\ldots,i_d}\}$ indexed by multi-indices, with overlap matrices computed along grid edges. Such a grid is literally a finite cubical complex $\mathscr{G}$, and the overlap data defines a discrete connection on it. The algorithm's holonomy check—computing $W_\gamma = \prod_{\text{edges}} U$ around each elementary plaquette and testing $\|W_\gamma - I\| < \varepsilon$—is the discrete box-filling condition of Corollary 5.4: can a parallel section defined on three faces of a 2-cube be extended to the fourth?

The cubical framework is therefore not an analogy imposed from outside. The computational objects are cubical sets, and the algorithm checks the fibration condition. Classical differential geometry provides the continuous theory ($F = 0$ on smooth manifolds), but the computation happens on a finite cubical complex where the relevant notions are combinatorial, not smooth. The adiabatic transport presheaf $\mathscr{A}$ restricted to $\mathscr{G} \hookrightarrow S_\square(M)$ is the computational data itself.

As the grid spacing $h \to 0$, $\|W_\gamma - I\| \sim O(h^2)$: the discrete presheaf approaches a fibration, and the per-plaquette deviation provides a principled adaptive mesh criterion near conical intersections.

8 Quasi-Diabatization as Nearest Fibration

When $F \neq 0$, exact diabatization is impossible (Theorem 2.3). Practitioners use various "quasi-diabatic" representations that approximately minimize derivative coupling [15]. The variational formulation is: find the discrete connection closest to satisfying the cubical fibration condition.

Given a grid $\mathscr{G} = (V, \mathscr{E})$ with overlap matrices $U_{kk'} \in \mathrm{U}(N)$, seek unitary transformations $S_k \in \mathrm{U}(N)$ at each vertex minimizing the *cubical defect*:

$$\Delta(\{S_k\}) = \sum_{(k,k') \in \mathscr{E}} \|S_k^\dagger U_{kk'} S_{k'} - I_N\|_F^2. \tag{10}$$

The transformed overlap $U_{kk'}^{\mathrm{d}} = S_k^\dagger U_{kk'} S_{k'}$ is the discrete connection in the candidate diabatic basis.

The cost function is not ad hoc: by Theorem 5.2, $\min_{\{S_k\}} \Delta = 0$ on a simply connected grid iff the restricted presheaf $\mathscr{A}|_{\mathscr{G}}$ is fibrant. Minimizing Δ is seeking the *nearest fibrant object* to $\mathscr{A}$—the model-categorical content that distinguishes this from classical group synchronization (same optimization, no fibration interpretation).

Proposition 8.1 (Topological obstruction to quasi-diabatization)**.** *Let Σ be a simply connected region in $\mathscr{G}$ with $\nu(\Sigma) \neq 0$. Then $\Delta^* := \min_{\{S_k\}} \Delta > 0$.*

Proof. If $\Delta^* = 0$, then $U_{kk'}^{\mathrm{d}} = I$ for all edges, so $W_f^{\mathrm{d}} = I$ for all plaquettes f, giving $\nu(\Sigma) = 0$—a contradiction. □

On a simply connected grid, the following are equivalent: $\min_{\{S_k\}} \Delta = 0$; all plaquette holonomies $W_f = I$; the ASOA adds no formal fillers (Section 5.5); and $\nu(\Sigma) = 0$ for every simply connected Σ. The cubical defect thus measures the same obstruction the ASOA resolves formally: minimizing Δ seeks the gauge in which the fewest ASOA fillers are needed, while ν certifies when this obstruction cannot be eliminated. The chemistry consequence: no quasi-diabatization method can achieve zero residual coupling on a region with $\nu \neq 0$, and unlike existing methods [3, 4, 16], the cubical defect is gauge-invariant by construction and subject to this topological lower bound.

Remark 8.2 (Scope of the variational principle)**.** *When overlap matrices are $U_{kk'} = V_k^\dagger V_{k'}$ with square unitary $V_k \in \mathrm{U}(N)$, the gauge $S_k = V_k^\dagger V_{k_0}$ gives $S_k^\dagger U_{kk'} S_{k'} = I$ on every edge, so $\Delta = 0$ identically. This does not invalidate the variational principle; it identifies its scope. The cubical defect is non-trivial for:* (a) *abstract discrete* $\mathrm{U}(N)$ *connections on graphs (the paper's formal setting);* (b) subspace tracking, *where $V_k \in \mathbb{C}^{M \times N}$ with $M > N$ is rectangular and $V_k V_k^\dagger \neq I_M$—the standard multi-reference chemistry setting ($N = 2$–10 states tracked within dimension $M \gg N$). The topological lower bound* (*Proposition* 8.1) *applies in both cases; Appendix* F *verifies both predictions on a 4-state model.*

9 Conclusion

The adiabatic transport presheaf (Definition 2.2) provides a cubical set construction whose fibration property translates, via the Frobenius theorem, to vanishing of Berry curvature (Theorem 2.3). The cubical face map identities encode holonomy (Theorem 5.2), ruling out smooth parallel-transport-based fibrant replacement when $F \neq 0$ (Corollary 5.5).

These results yield four positive outputs. First, the ASOA fibrant replacement (Corollary 5.7) canonically recovers spanning-tree diabatization from the model structure, with no chemical input and one canonical resolution per open box type. Second, discrete Berry numbers $\nu \in (\mathbb{Z}/2)^{N-1}$ (Theorem 6.2) classify CI configurations from grid overlap data; this is validated on a model Hamiltonian (Appendix C) and on the D_{3h} conical intersection in H_3 from ab initio SA-CASSCF wavefunctions (Appendix G). Third, quasi-diabatization (Section 8) is formulated as seeking the nearest fibrant object, with a topological lower bound from ν; its scope is identified in Remark 8.2 and verified numerically in the subspace-tracking regime (Appendix F). Fourth, the space of diabatization schemes is the flag manifold $\mathrm{Fl}(\mathbb{C}^N) = \mathrm{U}(N)/\mathrm{U}(1)^N$ (Theorem 3.2), with $\pi_2 = \mathbb{Z}^{N-1}$ detecting multi-state topological obstructions beyond the classical Berry phase.

Appendix D presents a Cubical Agda formalization certifying the algebraic foundations for $\mathbb{Z}/2$-valued connections, and discusses the relationship to CCHM cubical type theory [17]. A sharp limit: the fibers $E_x \cong \mathbb{C}^N$ are contractible, so $\mathrm{Fl}(\mathbb{C}^N)$ cannot be recovered from Awodey's classifying type—the cubical framework captures connection-level structure, while gauge orbit classification enters through differential geometry (Theorem 3.2).

References

[1] M. V. Berry. Quantal phase factors accompanying adiabatic changes. *Proc. R. Soc. London A*, 392:45–57, 1984. doi:10.1098/rspa.1984.0023

[2] B. Simon. Holonomy, the quantum adiabatic theorem, and Berry's phase. *Phys. Rev. Lett.*, 51:2167–2170, 1983. doi:10.1103/PhysRevLett.51.2167

[3] S. F. Boys. Construction of some molecular orbitals to be approximately invariant for changes from one molecule to another. *Rev. Mod. Phys.*, 32:296–299, 1960. doi:10.1103/RevModPhys.32.296

[4] T. Pacher, L. S. Cederbaum, and H. Köppel. Adiabatic and quasidiabatic states in a gauge theoretical framework. *Adv. Chem. Phys.*, 84:293–391, 1993. doi:10.1002/9780470141427.ch4

[5] D. R. Yarkony. Diabolical conical intersections. *Rev. Mod. Phys.*, 68:985–1013, 1996. doi:10.1103/RevModPhys.68.985

[6] R. J. Cave and M. D. Newton. Generalization of the Mulliken–Hush treatment for the calculation of electron transfer matrix elements. *Chem. Phys. Lett.*, 249:15–19, 1996. doi:10.1016/0009-2614(95)01310-5

[7] F. Wilczek and A. Zee. Appearance of gauge structure in simple dynamical systems. *Phys. Rev. Lett.*, 52:2111–2114, 1984. doi:10.1103/PhysRevLett.52.2111

[8] H. C. Longuet-Higgins, U. Öpik, M. H. L. Pryce, and R. A. Sack. Studies of the Jahn–Teller effect. II. The dynamical problem. *Proc. R. Soc. London A*, 244:1–16, 1958. doi:10.1098/rspa.1958.0022

[9] S. Awodey. *Cartesian Cubical Model Categories*. Lecture Notes in Mathematics 2385, Springer, 2025. `arXiv:2305.00893`. doi:10.1007/978-3-032-08730-0

[10] S. Awodey, E. Cavallo, T. Coquand, E. Riehl, and C. Sattler. The equivariant model structure on cubical sets. `arXiv:2406.18497`, 2024.

[11] A. Bohm, A. Mostafazadeh, H. Koizumi, Q. Niu, and J. Zwanziger. *The Geometric Phase in Quantum Systems*. Springer, 2003. doi:10.1007/978-3-662-10333-3

[12] C. A. Mead and D. G. Truhlar. Conditions for the definition of a strictly diabatic electronic basis for molecular systems. *J. Chem. Phys.*, 77:6090–6098, 1982. doi:10.1063/1.443853

[13] R. Garner. Understanding the small object argument. *Appl. Categ. Structures*, 17(3):247–285, 2009. doi:10.1007/s10485-008-9137-4

[14] E. Riehl. *Categorical Homotopy Theory*. New Mathematical Monographs 24, Cambridge University Press, 2014. doi:10.1017/CBO9781107261457

[15] W. Domcke, D. R. Yarkony, and H. Köppel, editors. *Conical Intersections: Electronic Structure, Dynamics and Spectroscopy*. World Scientific, 2004. doi:10.1142/5406

[16] K. R. Yang, X. Xu, and D. G. Truhlar. Direct diabatization of electronic states by the fourfold way. *Chem. Phys. Lett.*, 573:84–89, 2013. doi:10.1016/j.cplett.2013.04.036

[17] C. Cohen, T. Coquand, S. Huber, and A. Mörtberg. Cubical type theory: A constructive interpretation of the univalence axiom. *J. Automated Reasoning*, 60:391–455, 2018. doi:10.1007/s10817-018-9472-6

[18] A. Vezzosi, A. Mörtberg, and A. Abel. Cubical Agda: A dependently typed programming language with univalence and higher inductive types. *J. Funct. Program.*, 31:e8, 2021. doi:10.1017/S0956796821000034

[19] M. Bezem, T. Coquand, and S. Huber. The univalence axiom in cubical sets. *J. Eur. Math. Soc.*, 21:3343–3373, 2019. doi:10.4171/JEMS/905

[20] U. Schreiber. Differential cohomology in a cohesive infinity-topos. `arXiv:1310.7930`, 2013.

[21] Q. Sun, X. Zhang, S. Banerjee, P. Bao, M. Barbry, N. S. Blunt, *et al.* Recent developments in the PySCF program package. *J. Chem. Phys.*, 153:024109, 2020. doi:10.1063/5.0006074

A The Diabatization Problem: Context for Two Audiences

This appendix provides context for readers from either community: category theorists unfamiliar with quantum chemistry, and chemists unfamiliar with cubical homotopy theory. It can be read independently of the main text.

Why diabatization matters. Many processes in photochemistry—vision, photosynthesis, solar energy conversion, radiation damage to DNA—involve molecules passing through *conical intersections*, points in nuclear configuration space where two electronic energy surfaces become degenerate. At these points, the Born-Oppenheimer approximation breaks down and electrons can transition between energy surfaces on femtosecond timescales, making non-adiabatic transitions the rule rather than the exception in photochemistry. Simulating these processes requires solving the nuclear Schrödinger equation on multiple coupled potential energy surfaces simultaneously. In the standard (adiabatic) representation, the derivative couplings between surfaces diverge as $1/R$ at conical intersections, making numerical integration ill-conditioned or impossible. The *diabatic representation* replaces these singular couplings with smooth off-diagonal potential energy terms, enabling stable numerical propagation. Finding such a representation—the *diabatization problem*—is a prerequisite for essentially all multi-state quantum dynamics simulations, including trajectory surface hopping, multi-configurational time-dependent Hartree, and variational multi-configurational Gaussian methods.

The problem, for mathematicians. In quantum chemistry, the Born-Oppenheimer approximation produces electronic states $|\psi_i(R)\rangle$ that depend parametrically on the nuclear configuration $R \in M$. These states form a rank-N Hermitian vector bundle $E \to M$ with a natural $\mathrm{U}(N)$ connection (the Berry connection). The states are smooth away from *conical intersections*—codimension-2 submanifolds where eigenvalues become degenerate and the eigenvector assignment is discontinuous.

Simulating nuclear dynamics requires the electronic Hamiltonian as a smooth matrix-valued function of R. In the eigenbasis (adiabatic representation), this matrix is diagonal but has singular derivative couplings $\langle\psi_i|\nabla_R|\psi_j\rangle$ near conical intersections. A *diabatic representation* is a different basis—a gauge transformation—in which the Hamiltonian is non-diagonal but smooth. Finding such a basis is the diabatization problem: it is a gauge-fixing problem on the Berry connection.

On a simply connected region free of conical intersections, the Berry connection is flat and parallel transport gives an exact diabatic basis. When the region contains conical intersections (non-trivial holonomy), exact diabatization is impossible and one seeks an approximate solution—*quasi-diabatization*.

The standard computational method is: (i) compute eigenstates on a grid of nuclear configurations; (ii) form overlap matrices between adjacent grid points; (iii) propagate a reference frame along a spanning tree of the grid by composing these overlaps. This is discrete parallel transport on a cubical complex—precisely the setting of cubical homotopy theory.

The problem, for chemists. This paper recasts the diabatization algorithm in the language of *cubical sets* and *fibrations*. The dictionary is:

Your computational grid = a finite cubical complex
Overlap matrices $U_{kk'}$ = a discrete $\mathrm{U}(N)$ connection
Holonomy check: $\|W_\gamma - I\| < \varepsilon$? = testing the cubical fibration condition
Spanning-tree propagation = fibrant replacement (ASOA)
CI enclosed $\Rightarrow$ inexact diabatization = topological obstruction to fibration

The categorical framework provides three outputs beyond repackaging. First, a discrete topological invariant $\nu \in (\mathbb{Z}/2)^{N-1}$ that classifies conical intersection types from grid data—extracting information the standard $\|W - I\| > \varepsilon$ check currently discards (Theorem 6.2). Second, a variational formulation of quasi-diabatization with a topological lower bound on residual coupling (Section 8). Third, a proof that the spanning-tree algorithm is not merely a heuristic but the *canonical* output of a general categorical construction (Corollary 5.7), providing theoretical justification for a procedure used since the 1990s.

What the Cubical Agda formalization certifies. The spanning-tree algorithm's correctness depends on three algebraic properties: gauge invariance of holonomy (the result doesn't depend on the input basis), additivity of Berry numbers (invariants compose when regions are joined), and the equivalence between flatness and fillability (parallel transport gives an exact diabatic basis iff the connection is flat). These properties have been machine-verified in Cubical Agda [18] for the $\mathbb{Z}/2$ Berry parity model (Appendix D). The type-checker certifies that these properties hold by construction in the discrete cubical framework, independently of the continuous theory they discretize.

Diabatization methods as points in $\mathrm{Fl}(\mathbb{C}^N)$**.** The flag manifold classification (Theorem 3.2) organizes the landscape of diabatization methods: each method selects a point or submanifold. The assignment respects bundle isomorphisms: a connection-preserving isomorphism induces a diffeomorphism of flag manifolds mapping each method's locus to itself.

Method	**Selection rule**	**Locus in $\mathrm{Fl}(\mathbb{C}^N)$**	**Unique?**
Parallel transport	Path-ind. lift from R_0	Single point	Yes
Boys localization [3]	$\max \sum \lvert\mu_{ii}\rvert^2$	Critical submanifold	Locally
Block diag. [4]	Eigenbasis of P	Eigenspace flag	If non-deg.
Gen. Mulliken-Hush [6]	Eigenbasis of $\hat{\mu}$	Dipole-eigenspace flag	If non-deg.

Parallel transport from a basepoint R_0 is the geometric analogue of fibrant replacement. Boys localization maximizes $\sum_i |\langle \phi_i | \hat{\mu} | \phi_i \rangle|^2$ (the dipole operator $\hat{\mu}$), with critical points forming a submanifold of $\mathrm{Fl}(\mathbb{C}^N)$.

Block diagonalization diagonalizes a molecular property matrix P; k-fold degeneracy leaves $\mathrm{U}(k)/\mathrm{U}(1)^k$ residual freedom. The geodesic distance on $\mathrm{Fl}(\mathbb{C}^N)$ from the bi-invariant metric on $\mathrm{U}(N)$ provides a natural measure of disagreement between methods.

B The Diabatization Algorithm

The parallel transport construction (the geometric fibrant replacement) yields a computational algorithm, stated here for completeness.

Given adiabatic states $\{|\psi_i(R_k)\rangle\}$ on a grid $\mathscr{G} = (V, \mathscr{E})$ of nuclear configurations:

Step 1 (Transport matrices). For each edge $(R_k, R_{k'})$, compute the overlap $S = \Psi_k^\dagger \Psi_{k'}$ and extract the unitary factor via SVD: $S = U\Sigma V^\dagger$, $U_{k\to k'} = UV^\dagger$.

Step 2 (Topology detection). For each elementary cycle γ in $\mathscr{G}$, compute the holonomy $W_\gamma = \prod_{\text{edges}} U$. If $\|W_\gamma - I\| > \varepsilon$, the cycle encloses a conical intersection.

Step 3 (Diabatic basis). Fix a basepoint k_0 with $S_{k_0} = I_N$. For each vertex R_k in BFS order from k_0, set $S_k = (\prod_{\text{path edges}} U)^\dagger$. The diabatic Hamiltonian is $H_k^{\text{diab}} = S_k H_k^{\text{adiab}} S_k^\dagger$.

This algorithm is the discrete analogue of fibrant replacement: the BFS spanning tree provides canonical paths from the basepoint, and parallel transport along these paths “flattens” the connection. On simply connected domains the result is independent of spanning tree choice (up to numerical precision).

C Three-State Model Computation

We verify the discrete CI classification (Theorem 6.2) on a concrete three-state model Hamiltonian with two isolated conical intersections.

C.1 Model Hamiltonian

Consider $M = \mathbb{R}^2$ with coordinates (x, y) and the 3×3 diabatic Hamiltonian

$$H^{\mathrm{d}}(x,y) = \begin{pmatrix} x+a & y & 0 \\ y & 0 & y \\ 0 & y & -(x-a) \end{pmatrix} \tag{11}$$

with $a > 0$ a separation parameter. This has two conical intersections: $\mathrm{CI}_{1\text{-}2}$ at $(-a, 0)$ where $E_1 = E_2$ and the $(1,2)$-coupling vanishes, and $\mathrm{CI}_{2\text{-}3}$ at $(a, 0)$ where $E_2 = E_3$ and the $(2,3)$-coupling vanishes. Near each CI, the local structure is a standard Jahn–Teller cone.

C.2 Discrete Computation on a Grid

Place a rectangular grid $\mathscr{G}$ on $[-L, L]^2$ with spacing h. At each grid point R_k, diagonalize $H^{\mathrm{d}}(R_k)$ to obtain adiabatic states $\{|\psi_i(R_k)\rangle\}_{i=1}^3$. For each edge $(R_k, R_{k'})$, compute the 3×3 overlap matrix $U_{kk'} = \Psi_k^\dagger \Psi_{k'}$.

For three loops in $\mathscr{G}$:

- γ_{12}: rectangle enclosing only $(-a, 0)$
- γ_{23}: rectangle enclosing only $(a, 0)$
- γ_{both}: rectangle enclosing both CIs

compute the holonomy $W_\gamma = \prod_{e\in\gamma} U_e$ and extract eigenvalues $e^{i\gamma_1}, e^{i\gamma_2}, e^{i\gamma_3}$ ordered by continuity from the adiabatic basis at the basepoint. The multi-state Berry numbers $\nu = (\nu_1, \nu_2)$ are computed as in Definition 6.1.

For $a = 2$, $L = 4$, $h = 0.1$, we obtain:

Loop	(b_1, b_2, b_3)	ν
γ_{12}	$(1,1,0)$	$(1,0)$
γ_{23}	$(0,1,1)$	$(0,1)$
γ_{both}	$(1,0,1)$	$(1,1)$

matching the predictions of Section 4 and confirming the additivity $\nu(\gamma_{\text{both}}) = \nu(\gamma_{12}) + \nu(\gamma_{23})$ in $(\mathbb{Z}/2)^2$. The scalar holonomy check $\|W - I\|$ detects non-triviality for all three loops, but cannot distinguish γ_{both} from a single CI: all three have eigenvalue multiset $\{-1,-1,+1\}$. The $(\mathbb{Z}/2)^2$ invariant resolves this ambiguity completely.

C.3 Comparison with Continuous Computation

In the smooth setting, the same information requires: (i) constructing a smooth Berry connection A by interpolating from grid data, (ii) choosing a smooth surface Σ with $\partial\Sigma = \gamma$, and (iii) integrating $F = dA + A \wedge A$ over Σ to obtain Chern numbers. For the discrete computation, none of these steps is needed: the overlap matrices are the primary data, the grid plaquettes constitute the surface, and the Berry numbers are extracted from the eigenvalues of the holonomy product. The additional computational cost beyond the standard algorithm is one $N \times N$ eigenvalue decomposition per loop.

D Certified Computation in Cubical Agda

We have formalized the discrete Berry number computation in Cubical Agda [18], producing machine-checked proofs of the core algebraic properties. The formalization works over $\mathbb{Z}/2$-valued connections—the Berry parity model—which captures the essential discrete structure while remaining tractable for formal verification. Source code is available at `https://github.com/karsar/berry_agda`.

Module structure. (i) `Z2`: $\mathbb{Z}/2$ as a group under XOR, with certified associativity, commutativity, identity, and involution. (ii) `Grid`: finite cubical grids with vertices, directed edges, plaquettes, $\mathbb{Z}/2$-valued connections, holonomy as XOR around plaquettes, and gauge transformations. (iii) `GaugeInvariance`: proof that holonomy is invariant under gauge transformation, verified by exhaustive case analysis (256 cases). (iv) `BerryNumber`: Berry number defined as plaquette holonomy, with additivity for adjacent plaquettes sharing an edge: shared edge values cancel by symmetry plus an algebraic identity (128 cases). (v) `Fibration`: the open box filler (unique fourth edge achieving zero holonomy), with proofs that the filler is correct and that flatness is equivalent to fillability (discrete Theorem 2.3).

Certified properties. Three properties from the main text are formalized as types whose inhabitants are constructive proofs:

- *Gauge invariance* (Theorem 6.2(a)): `gauge-invariant` proves $\mathrm{hol}(g \cdot A, f) = \mathrm{hol}(A, f)$ for all gauges g, connections A, and plaquettes f.

- *Additivity* (Theorem 6.2(b)): `berry-additivity` proves $\nu(f_1) \oplus \nu(f_2) = \nu(f_1 \cup f_2)$ for adjacent plaquettes with symmetric connections.
- *Fibration ⇔ flatness* (Theorem 2.3): `filler-correct` and `flat-fills` together prove that the open box condition is equivalent to zero holonomy.

Scope and significance. The formalization covers the $\mathbb{Z}/2$ Berry parity model, not full $\mathrm{U}(N)$ connections. An intermediate extension to $\mathbb{Z}$-valued connections—tracking integer Chern numbers, with holonomy as addition around plaquettes—would capture the full $\mathbb{Z}^{N-1}$ classification without requiring eigenvalue decomposition, and is a tractable next step. Full $\mathrm{U}(N)$ would require formalized matrix algebra and eigenvalue extraction, a substantially larger project. Nevertheless, the $\mathbb{Z}/2$ model captures the algebraic structure that underlies gauge invariance, additivity, and the fibration–flatness equivalence. The formalization provides independent verification that these properties hold by construction in the discrete cubical framework, not merely by appeal to the continuous theory they discretize. More broadly, it demonstrates that the cubical framework admits formal verification: the type-checker certifies ν, guaranteed correct by construction rather than by floating-point comparison.

Cartesian versus De Morgan cubical structure. The theoretical framework of this paper (Sections 2–8) uses Cartesian cubical sets in the sense of Awodey [9]: face maps, degeneracies, and open box filling, with no reversal operation on the interval. The Cubical Agda implementation, however, uses CCHM-style cubical type theory [18], which equips the interval with De Morgan structure: a reversal involution $\sim i = 1 - i$ that provides path reversal as a primitive. For the $\mathbb{Z}/2$ Berry parity model, this is a natural fit: every element of $\mathbb{Z}/2$ is its own inverse, so reversing a path preserves holonomy ($\mathrm{hol}(\gamma^{-1}) = \mathrm{hol}(\gamma)$ for $\mathbb{Z}/2$-valued connections). The De Morgan reversal encodes this involution directly as a structural identity, rather than requiring a separate proof. This distinction has no effect on the certified properties (gauge invariance, additivity, and fibration–flatness hold in both settings), but the De Morgan structure simplifies the formalization of path reversal arguments that arise in the gauge invariance proof.

Relationship to cubical type theory. In the CCHM semantics of cubical type theory [17], a fibration $p : E \to B$ models a dependent type whose Kan filling operations give transport along paths. Theorem 2.3 says that the adiabatic transport presheaf models such a dependent type if and only if $F = 0$—filling operations exist iff the connection is flat. When $F \neq 0$, the dependent type of simultaneously parallel sections has no global elements over 2-cubes (Theorem 5.2), and the ASOA produces formal fillers that witness a non-geometric resolution. The fibers $E_x \cong \mathbb{C}^N$ are contractible, so the presheaf carries no higher homotopical information in the sense of Bezem-Coquand-Huber [19] or Awodey [9]; the content is entirely at the level of connections and transport, not homotopy types.

Type-theoretic reading of the main results. The main text uses Awodey's Cartesian cubical framework [9], whose model-categorical properties (Theorem 9.8, op. cit.) underpin the ASOA analysis. However, the CCHM semantics provides a complementary reading of each main theorem in type-theoretic language, which we record here for readers from the cubical type theory community.

The Berry connection provides *transport without composition*. In CCHM, a fibration has both transport (lifting along 1-paths) and composition (coherent filling of 2-cubes). Theorem 2.3(1) says the adiabatic transport presheaf always has 1-transport (ODE existence), while parts (2)–(3) say composition exists iff $F = 0$. Thus the Berry connection is a *pre-type family with transport but without composition*—an object that arises naturally in gauge theory but is rarely encountered in type theory, where fibration conditions are all-or-nothing.

The holonomy obstruction (Theorem 5.2) has a precise compositional reading. The CCHM `comp` operation takes a partial tube (open box) and produces a filler; the cubical identity $\partial_1^1\partial_2^1 = \partial_1^1\partial_1^1$ on the transport candidate is the statement that composing transport left→top equals transport bottom→right—i.e., path composition is coherent. Its failure is exactly holonomy. The ASOA fillers (Proposition 5.6) are then freely postulated `comp` terms adjoined to a pre-type family that lacks them: passing from $\mathscr{A}$ to $\hat{\mathscr{A}}$ freely adds composition witnesses not arising from geometric transport.

Finally, the contractible-fiber limitation explains why gauge structure is invisible to the homotopy-theoretic layer. In CCHM, a fibration with contractible fibers is equivalent to the identity family (all types in the family are propositionally equal to $\mathbf{1}$). Since our fibers $E_x \cong \mathbb{C}^N$ are contractible, $\hat{\mathscr{A}}$ is homotopically trivial—confirming the sharp limit stated in Section 9. All gauge-theoretic content (holonomy, $\mathrm{Fl}(\mathbb{C}^N)$ classification, ν invariants) resides in the *connection structure* that the face maps encode, not in the homotopy type. Recovering this connection-level information type-theoretically would require enriching the type theory with differential structure, as in the cohesive framework of Schreiber [20]; we leave this direction to future work.

E Explicit ASOA on a Multi-Plaquette Grid

We compute the ASOA fibrant replacement on a 2×1 grid with a 3-state system, demonstrating the transport formula of Proposition 5.6 on explicit matrices and establishing the scope of the ASOA's contribution.

Setup. The grid has vertices $(i, j) \in \{0, 1, 2\} \times \{0, 1\}$, two plaquettes f_1 (corner $(0,0)$) and f_2 (corner $(1,0)$), and seven directed edges. Set all bottom, left, and right-boundary transports to I_3. The shared edge (vertical, from $(1,0)$ to $(1,1)$) carries a generic unitary $U_e = \mathrm{diag}(e^{i\alpha}, e^{i\beta}, e^{i\gamma})$. Holonomies $W_f = \tau_L^{-1}\tau_T^{-1}\tau_R\tau_B$ then reduce to

$$W_{f_1} = \tau_{T_1}^{-1}U_e, \qquad W_{f_2} = U_e^{-1}\tau_{T_2}^{-1}, \tag{12}$$

where τ_{T_k} is the top-edge transport of plaquette f_k. Choose

$$W_{f_1} = \mathrm{diag}(-1, -1, 1), \qquad W_{f_2} = \mathrm{diag}(1, -1, -1), \tag{13}$$

a $(1,2)$-CI in f_1 and a $(2,3)$-CI in f_2, with Berry numbers $\nu(f_1) = (1,0)$ and $\nu(f_2) = (0,1)$. The top transports are then $\tau_{T_1} = U_eW_{f_1}^{-1}$ and $\tau_{T_2} = W_{f_2}^{-1}U_e^{-1}$ (which equal $U_eW_{f_1}$ and $W_{f_2}U_e^{-1}$ respectively, since the chosen holonomies are involutory: $W_{f_k}^2 = I$).

ASOA Step 1. Proposition 5.6 gives the missing-top filler $\hat{\tau}_T = \tau_R\tau_B\tau_L^{-1}$. In our gauge:

$$\hat{\tau}_{T_1} = U_e \cdot I \cdot I = U_e, \tag{14}$$

$$\hat{\tau}_{T_2} = I \cdot I \cdot U_e^{-1} = U_e^{-1}. \tag{15}$$

These formal transports differ from the geometric ones by the holonomy:

$$\hat{\tau}_{T_1} = \tau_{T_1} \cdot W_{f_1}, \qquad \hat{\tau}_{T_2} = \tau_{T_2} \cdot W_{f_2}, \tag{16}$$

confirming $\hat{\tau}_T = \tau_T \cdot \mathrm{Ad}_{\tau_L}(W_f)$ from (8) (with $\tau_L = I$ for f_1 and $\tau_L = U_e$ for f_2; note $\mathrm{Ad}_{U_e}(W_{f_2}) = W_{f_2}$ since both are diagonal).

For each plaquette the missing-right filler is $\hat{\tau}_R = \tau_T \tau_L \tau_B^{-1}$:

$$\hat{\tau}_R^{f_1} = \tau_{T_1}, \qquad \hat{\tau}_L^{f_2} = \tau_{T_2}^{-1}, \tag{17}$$

where $\hat{\tau}_L^{f_2}$ uses the missing-left formula $\hat{\tau}_L = \tau_T^{-1}\tau_R\tau_B$. Both act on the shared edge.

Step 2 and termination. On a 2D grid, there are no 3-cubes, so Step 1 fillers create no new open 2-boxes requiring further filling. The ASOA terminates at Step 1: $\hat{\mathscr{A}} = \hat{\mathscr{A}}_1$.

At the shared edge, the Step 1 fillers from f_1 and f_2 provide two formal transports, $\hat{\tau}_R^{f_1} = \tau_{T_1}$ and $\hat{\tau}_L^{f_2} = \tau_{T_2}^{-1}$, alongside the geometric U_e. Their discrepancy is

$$(\hat{\tau}_R^{f_1})^{-1}\hat{\tau}_L^{f_2} = \tau_{T_1}^{-1}\tau_{T_2}^{-1} = (W_{f_1}U_e^{-1})(U_e W_{f_2}) = W_{f_1}W_{f_2} = \mathrm{diag}(-1,1,-1), \tag{18}$$

which is entirely determined by W_{f_1} and W_{f_2}—the shared edge parameter U_e cancels.

Non-abelian variant. The diagonal example above has $\mathrm{Ad}_{U_e}(W_f) = W_f$ because all matrices commute. To exhibit the adjoint action's role, replace the shared edge transport by

$$U_e = R_{13}(\pi/4) \;=\; \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 0 & -1 \\ 0 & \sqrt{2} & 0 \\ 1 & 0 & 1 \end{pmatrix}, \tag{19}$$

a rotation mixing states 1 and 3. This changes the top transports τ_{T_k} (so that holonomies remain (13)) but leaves W_{f_1}, W_{f_2}, and all boundary transports unchanged.

The Step 1 fillers are still $\hat{\tau}_{T_k} = \tau_{T_k} \cdot \mathrm{Ad}_{\tau_L^{(k)}}(W_{f_k})$. For f_1, $\tau_L^{(1)} = I$ and the correction is simply W_{f_1}, independent of U_e. For f_2, $\tau_L^{(2)} = U_e = R_{13}(\pi/4)$ and

$$\mathrm{Ad}_{R_{13}(\pi/4)}\big(\mathrm{diag}(1,-1,-1)\big) \;=\; \begin{pmatrix} 0 & 0 & 1 \\ 0 & -1 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \tag{20}$$

which has the same eigenvalues $\{1,-1,-1\}$ as W_{f_2} but is non-diagonal: the correction now mixes states 1 and 3. Compare with the abelian case, where $\mathrm{Ad}_I(W_{f_2}) = \mathrm{diag}(1,-1,-1)$ acts only on state 2.

Both configurations have identical plaquette holonomies, identical Berry numbers, and identical shared-edge discrepancy (18). But their ASOA fillers differ:

$$\hat{\tau}_{T_2}^{(\mathrm{diag})} = \tau_{T_2}^{(\mathrm{diag})} \cdot \mathrm{diag}(1,-1,-1), \qquad \hat{\tau}_{T_2}^{(R_{13})} = \tau_{T_2}^{(R_{13})} \cdot \begin{pmatrix} 0 & 0 & 1 \\ 0 & -1 & 0 \\ 1 & 0 & 0 \end{pmatrix}. \tag{21}$$

The ASOA correction at f_2 records how f_1's Berry curvature defect appears when transported through the shared edge to f_2's gauge frame. When U_e mixes states 1 and 3, the $(1,2)$-CI defect at f_1 propagates as a correction mixing states 1 and 3 at f_2—precisely the non-abelian "holonomy defect propagation" of Section 5.5. This gauge-frame information is invisible to any gauge-invariant quantity (holonomies, Berry numbers, or the discrepancy (18)) but is retained by the ASOA as part of its canonical frame-adapted resolution.

Gauge-covariant structure. The preceding example illustrates a general principle: the ASOA operates in the *original gauge* (the adiabatic representation), producing fillers that are gauge-covariant rather than gauge-invariant. A spanning tree diabatization first transforms to tree gauge (setting all tree edges to I), then places residual holonomy on cut edges—destroying the original gauge frame. The ASOA instead distributes the holonomy defect via $\mathrm{Ad}_{\tau_L}(W_f)$ at each plaquette, preserving the relationship between the correction and the local adiabatic basis.

The ASOA correction $C_f = \mathrm{Ad}_{\tau_L}(W_f)$ at plaquette f satisfies:

1. C_f has the same eigenvalues as W_f (topology preserved);
2. C_f is conjugated into the gauge frame at the left edge of f (frame adapted to local transport);
3. for adjacent plaquettes sharing an edge, the corrections interact through the shared transport: in the example above, the correction at f_2 depends on U_e even though W_{f_2} does not, because Ad_{τ_L} conjugates the holonomy into the local frame.

A spanning tree diabatization achieves property (1) and fixes property (2) for tree edges, but breaks frame-adaptation at cut edges (where residual holonomy is concentrated), and different tree choices produce incompatible frames. The ASOA maintains all three properties by distributing the correction across every plaquette. Note, however, that when the ASOA fillers are converted to a vertex gauge $\{S_k\}$ by propagation from a root, the result is a specific spanning tree (Corollary 5.7): the gauge-covariant structure lives at the level of fillers, not of the extracted gauge.

Comparison with spanning trees. A spanning tree of the 2×1 grid cuts two edges (one per cycle). In tree gauge (all tree edges set to I), each cut edge carries a residual that absorbs the enclosed holonomy. Three representative trees, with residuals for the diagonal case (13):

Cut edges	*Residuals (on cut edges)*	*Label*
T_1, T_2	W_{f_1}, W_{f_2}	T_{top}
T_1, E	$W_{f_1}W_{f_2}, W_{f_2}$	T_{share}
E, T_2	$W_{f_1}, W_{f_1}W_{f_2}$	T'_{share}

All three are gauge-equivalent: the transformation $T_{\text{top}} \to T_{\text{share}}$ acts by $g_{(1,1)} = W_{f_2}$, trivially elsewhere. The ASOA fillers (14)–(15) constitute a fourth specific representative: replacing each geometric top transport by its ASOA filler produces a flat connection in the original gauge, equivalent to any spanning tree choice via $g_{(1,1)} = U_e^{-1}$.

Theorem E.1 (ASOA data on simply connected 2D grids)**.** *Let $\mathscr{G}$ be a simply connected finite 2D grid with $\mathrm{U}(N)$ connection $\{U_e\}$. The gauge equivalence class of the ASOA fibrant replacement $\hat{\mathscr{A}}|_{\mathscr{G}}$ is completely determined by the plaquette holonomy set $\{W_f : f \in \mathscr{G}\}$.*

Proof. On a simply connected grid, any connection is determined up to gauge by its holonomies on plaquettes (the grid's 2-cells generate H_2). Since the ASOA filler at each plaquette is determined by the presheaf $\mathscr{A}|_{\mathscr{G}}$, gauge-equivalent inputs yield gauge-equivalent outputs. On a 2D grid the ASOA terminates at Step 1 (no 3-cubes), so $\hat{\mathscr{A}} = \hat{\mathscr{A}}_1$, which is determined by $\{W_f\}$ up to gauge. $\square$

What the ASOA uniquely provides. The theorem shows the ASOA does not extract new gauge-invariant information from a simply connected 2D grid beyond $\{W_f\}$. Its contribution is structural: canonicity (one resolution vs. exponentially many spanning trees, Corollary 5.7) and three additional properties demonstrated by the example above.

1. *Gauge covariance*: the ASOA operates in the original gauge, producing corrections (20) adapted to the local adiabatic frame. Two connections with identical $\{W_f\}$ but different inter-plaquette transports produce different ASOA corrections (21).
2. *Formal cell structure*: $\hat{\mathscr{A}}$ is a cubical set with both geometric and formal cells. The number and location of formal 2-cells—one per non-flat plaquette per open box type—is a grid-level summary of Berry curvature distribution, computable without choosing a gauge.

However, when the ASOA gauge is extracted by propagation from a root vertex, the result coincides with the spanning tree rooted there (Corollary 5.7): the gauge-covariant structure of the fillers does not survive gauge extraction. Whether it can be exploited differently (e.g., as weights in iterative quasi-diabatization) remains open. On non-simply connected grids (e.g., periodic boundary conditions on an $n_x \times n_y$ torus), the connection carries additional gauge-invariant data beyond $\{W_f\}$: the holonomy of non-contractible loops. The cycle space of a torus grid has $\dim = |\mathscr{E}| - |V| + 1 = n_x n_y + 1$, while plaquette boundaries span a codimension-2 subspace (the boundary of a closed surface is zero), leaving a quotient $\cong \mathbb{Z}^2$ generated by the two non-contractible cycles. Concretely, a spanning tree has $n_x n_y + 1$ cut edges: $n_x n_y - 1$ dual to plaquettes (carrying contractible holonomy W_f) and 2 dual to the generating non-contractible cycles (carrying monodromy). As a consequence, Theorem E.1 fails: different spanning trees can differ in how they distribute the non-contractible holonomy, yielding gauge-inequivalent diabatizations. The ASOA still produces one canonical filler per open box type, but the resulting gauge class depends on both $\{W_f\}$ and the monodromy of the non-contractible cycles—strictly more data than in the simply connected case. A systematic treatment of this regime is relevant for periodic systems in solid-state physics, and we leave it to future work.

F Subspace Tracking: Numerical Verification

Remark 8.2 identifies two regimes for the cubical defect (10): trivially zero when the full eigenspace is tracked ($N = M$, square unitary V_k), and non-trivially positive for subspace tracking ($N < M$, rectangular V_k). We verify both predictions on the following model.

Model. A four-state diabatic Hamiltonian on $\mathbb{R}^2$ with a conical intersection between states 1–2 at $(-a, 0)$:

$$H(x,y) = \begin{pmatrix} x+a & y & c & 0 \\ y & -(x+a) & 0 & c/2 \\ c & 0 & \Delta_g+0.1x & 0.05y \\ 0 & c/2 & 0.05y & \Delta_g+0.3-0.1x \end{pmatrix},$$

where $a = 2$, $\Delta_g = 5$ is the energy gap to the untracked states, and c is the inter-subspace coupling. When $c = 0$ the upper 2×2 block has a standard Jahn–Teller CI; for $c > 0$ the tracked and untracked subspaces mix, and the CI shifts slightly but persists (codimension-2 degeneracies are generically stable under perturbation).

We compute on a grid $[-a-L, -a+L] \times [-L, L]$ with $L = 2$, spacing $h = 0.5$ ($9 \times 9 = 81$ vertices). At each vertex k we diagonalize $H(x_k, y_k)$ to obtain eigenvectors $V_k \in \mathbb{R}^{4 \times N_{\text{track}}}$, then form overlap matrices $U_{kk'} = U_{\text{polar}}(V_k^\dagger V_{k'})$, where $U_{\text{polar}}(A) = UV^\dagger$ for the SVD $A = U\Sigma V^\dagger$ is the nearest unitary to A in Frobenius norm (for square unitary V_k, this equals $V_k^\dagger V_{k'}$ directly).

Factorization theorem. For the full eigenspace ($N = M = 4$), the factorization gauge $S_k = V_k^{\dagger} V_{k_0}$ gives $S_k^{\dagger} U_{kk'} S_{k'} = I$ on every edge, hence $\Delta_{\text{fac}} = 0$ to machine precision. For the tracked subspace ($N = 2$, $M = 4$), $V_k \in \mathbb{R}^{4\times 2}$ satisfies $V_k V_k^{\dagger} \neq I_4$, and the factorization breaks: $\Delta_{\text{fac}} > 0$.

Table 1 confirms both predictions across a range of coupling strengths.

Table 1: Cubical defect vs. inter-subspace coupling c. $\Delta_{\text{fac}}(N{=}4)$ is identically zero (machine precision) for all couplings, confirming the factorization theorem. $\Delta_{\text{fac}}(N{=}2)$ grows with c, measuring the failure of the rectangular overlap to factor. Δ_{tree} is the spanning-tree gauge value; Δ^* is the global optimum over $\mathrm{O}(2)$ gauges (81 rotation angles, 30 restarts of L-BFGS-B seeded from the spanning tree).

c	$\Delta_{\text{fac}}(N{=}4)$	$\Delta_{\text{fac}}(N{=}2)$	Δ_{tree}	Δ^*	$\Delta^*/\Delta_{\text{tree}}$
0.00	3×10^{-29}	0	0	0	—
0.05	2×10^{-28}	10^{-6}	$< 10^{-6}$	$< 10^{-7}$	0.107
0.10	2×10^{-28}	2×10^{-5}	$< 10^{-6}$	$< 10^{-7}$	0.107
0.20	5×10^{-28}	3.3×10^{-4}	3×10^{-6}	3×10^{-7}	0.108
0.30	3×10^{-28}	1.6×10^{-3}	1.4×10^{-5}	1.6×10^{-6}	0.108
0.50	5×10^{-28}	1.1×10^{-2}	9.6×10^{-5}	1.1×10^{-5}	0.110
0.80	5×10^{-28}	5.3×10^{-2}	4.5×10^{-4}	5.2×10^{-5}	0.115

Interpretation. At $c = 0$ the states decouple into 2×2 blocks and tracking $N = 2$ states is effectively a full-eigenspace problem within the block, so $\Delta = 0$. As coupling increases, the adiabatic eigenvectors at each grid point mix all four states, the tracked subspace projector $V_k V_k^{\dagger}$ rotates nontrivially, and the factorization identity $V_k V_k^{\dagger} = I$ that guarantees $\Delta = 0$ in Remark 8.2 fails progressively. The hierarchy $\Delta_{\text{fac}} \gg \Delta_{\text{tree}} \gg \Delta^* > 0$ persists across all nonzero couplings, with each step representing a different source of improvement: the factorization gauge $S_k = V_k^{\dagger} V_{k_0}$ accumulates distortion over long graph distances, while the spanning tree propagates locally; the global optimum Δ^* further improves by distributing residual coupling across the grid.

The ratio $\Delta^*/\Delta_{\text{tree}} \approx 0.108$ is stable across five decades of coupling strength, indicating that the improvement factor is geometric (determined by grid structure) rather than physical. The global optimum is approximately $9\times$ smaller than the spanning-tree value, but both track the same physical signal: the cubical defect introduced by inter-subspace mixing. Whether this ratio admits an analytic expression in terms of the grid dimensions and spanning-tree structure is an open question.

At strong coupling ($c = 0.8$), the loop holonomy eigenphases reach $\pm 0.52\pi$ and the discrete Berry number flips to $\nu = (1)$, activating the topological lower bound of Proposition 8.1: $\Delta^* = 5.2\times 10^{-5} > 0$, confirming that no gauge transformation can eliminate the cubical defect. (At intermediate coupling $c = 0.5$, the eigenphases are $\pm 0.46\pi$—below the $\pi/2$ rounding threshold for b_k—so the invariant reads $\nu = (0)$ despite nonzero holonomy. A finer grid enclosing the CI more tightly would shift the transition to smaller c.) This demonstrates that the invariant ν detects conical intersections through the subspace overlap connection even when the CI lies in a subsystem coupled to additional states.

G Ab Initio Validation: H_3 Conical Intersection

We validate the discrete Berry number (Theorem 6.2) on an ab initio electronic structure calculation, moving from model Hamiltonians (Appendix C) to a real chemical system.

System. Triatomic hydrogen H_3 has a conical intersection at the equilateral triangle (D_{3h}) geometry between its two lowest doublet states—the prototypical Berry phase system [8, 1, 12]. We compute two-state adiabatic wavefunctions at the SA-CASSCF(3,3)/STO-3G level (state-averaged complete active space SCF with 3 electrons in 3 orbitals, two equally weighted doublet states) using PySCF [21]. The third vertex of the triangle is displaced by Jahn–Teller coordinates (Q_x, Q_y) from the equilateral position; the CI lies at $Q_x = Q_y = 0$. The bond length is $r_0 = 1.8$ bohr (0.95 Å). At the D_{3h} geometry, the energy gap between the two states is zero to machine precision, confirming the degeneracy.

Method. At each pair of adjacent grid points or loop vertices $(R_k, R_{k'})$, we compute the 2×2 state overlap matrix $S^{IJ}_{kk'} = \langle \Psi_I(R_k) | \Psi_J(R_{k'}) \rangle$ via determinant-based CI overlap with cross-geometry AO integrals. Since SA-CASSCF produces real wavefunctions, the Berry parity for each state is determined by the sign of the diagonal overlap product around a closed loop: $b_j = 0$ if $\prod_k \langle \psi_j(R_k) | \psi_j(R_{k+1}) \rangle > 0$, and $b_j = 1$ otherwise. This is the per-state version of the holonomy eigenvalue test in Definition 6.1: for real wavefunctions, the eigenphases of W_γ are 0 or π, and the sign of the diagonal product determines which.

Circular loops. We compute the Berry parity around circular loops of radius 0.15 bohr centered at the CI, with varying numbers of sample points. All loops correctly yield $b = (1, 1)$, giving $\nu = 1$:

n_{pts}	$\prod_k \langle \psi_1^k \vert \psi_1^{k+1} \rangle$	$\prod_k \langle \psi_2^k \vert \psi_2^{k+1} \rangle$	b	ν
8	-0.519	-0.518	$(1,1)$	1
16	-0.724	-0.724	$(1,1)$	1
32	-0.852	-0.851	$(1,1)$	1
48	-0.899	-0.898	$(1,1)$	1

The overlap products are negative for all discretizations (converging toward -1 as the discrete approximation to the π-phase improves), correctly detecting the sign change predicted by Proposition 2.1. The result is equally stable across loop radii (0.05–0.30 bohr). For comparison, three off-CI loops (centers displaced to $(0.3, 0.3)$, $(-0.3, 0.3)$, $(0.0, 0.3)$ bohr, radius 0.10 bohr, 24 points each) all give $b = (0, 0)$, $\nu = 0$.

Grid plaquette map. On a 15×15 rectangular grid with $Q_x, Q_y \in [-0.25, 0.25]$ bohr (spacing $h = 0.036$ bohr), we compute the per-plaquette Berry number. Of the 196 plaquettes, exactly one has $\nu = 1$: the plaquette containing the origin, with center near $(Q_x, Q_y) = (0.018, 0.018)$ bohr. All other plaquettes have $\nu = 0$. The plaquette holonomy norm $\|W_f - I\|$ peaks at 2.82 on this cell (median over all cells: 2.8×10^{-5}), confirming that Berry curvature is sharply concentrated at the CI.

Concentric rectangular loops centered on the CI all give $\nu = 1$ regardless of loop size (margin 1–7 plaquettes), with per-state overlap products in the range $[-0.90, -0.53]$—consistently negative and well-separated from zero.

Significance. This is the complete pipeline envisioned by Theorem 6.2: ab initio wavefunctions $\to$ overlap matrices $\to$ holonomy products $\to$ Berry parity $\to$ discrete invariant $\nu \in \mathbb{Z}/2$. No smooth interpolation, surface integration, or Chern number computation is required. The invariant is computed entirely from the primary data of the standard diabatization algorithm (Appendix B), at the additional cost of tracking overlap signs around loops.