

# Coherence Types

A Refinement Calculus where Convention Coherence is the Declarations Forcing  
Every Reconciliation — Parts I (Statics) and II (Dynamics and Soundness)

John Komkov, Independent Researcher

June 2026

## Contents

Abstract	2
1 1. Introduction	3
1.1 1.1 The missing composition rule, as a typing problem	3
1.2 1.2 The thesis, stated precisely	3
1.3 1.3 Why a type system: the bridge from the hardness result	5
1.4 1.4 Contributions	5
1.5 1.5 Claim discipline and scope	5
2 2. The calculus $\lambda_{\nabla}$	6
2.1 2.1 Frames, declared transports, and opacity	6
2.2 2.2 Types	6
2.3 2.3 Terms	7
2.4 2.4 The convention system of a term; the context as presheaf	7
2.5 2.5 Two running examples	8
3 3. Statics	8
3.1 3.1 Constraint generation	8
3.2 3.2 Consistency, the fee, and decidability	9
3.3 3.3 The coherence judgment	10
3.4 3.4 Coercion is disclosure; minimum coercion is the fee	11
3.5 3.5 Frame polymorphism	12
3.6 3.6 The frame-rule reading	12
3.7 3.7 The boundaries	12
4 Part II. Dynamics and Soundness	13
5 4. Operational semantics	13
5.1 4.1 Runtime configurations	13
5.2 4.2 Reduction	13
6 5. Soundness	14

6.1	5.6 What Part II establishes; what is next . . . . .	15
7	Part III. The Graded Disclosure Modality and Interface Evolution	16
8	6. The graded disclosure modality $\square_r$	16
8.1	6.1 The grade monoid . . . . .	16
8.2	6.2 The modality and its comonad structure . . . . .	16
8.3	6.3 The two discharge operations (= [CD] §7.0) . . . . .	17
8.4	6.4 Graded typing; the grade is the fee . . . . .	18
8.5	6.5 Connection to graded / quantitative type theory . . . . .	18
8.6	6.6 Feedback and recursion, re-admitted . . . . .	18
8.7	6.7 Graded soundness: the grade is the repair cost . . . . .	18
9	7. Interface evolution as graded type preservation (discharging the abstract half of [CD] Conjecture 9.5)	19
9.1	7.1 The problem . . . . .	19
9.2	7.2 Typed interface updates . . . . .	20
9.3	7.3 The preservation theorem . . . . .	20
9.4	7.4 Stability: bounded recertification under general updates . . . . .	21
9.5	7.5 What Part III establishes . . . . .	21
10	References	22

## Abstract

Hoare’s rule for sequential composition is sound only when the predicate at the seam has a stable meaning on both sides. Open agent/ tool compositions break that assumption: each component satisfies its local contract, but the conventions the components attach to shared quantities — date bases, unit scales, currencies, timezones — are transported across interfaces and need not agree where a feedback loop closes. A companion program identifies the forced obstruction invariant — the coherence fee, the witness rank of a semantic interface complex (Komkov, A Witness Logic for Semantic Composition, 2026) — and a second companion proves it cannot be recovered from local observation (Komkov, Compositional Incoherence is a Parity, 2026: SQ-hard and LPN-hard). The negative result has a constructive corollary: if coherence cannot be inferred bottom-up, it must be carried top-down, as a typing invariant. This paper builds that type system, in two parts. Part I (statics) presents  $\lambda_\nabla$ , a refinement calculus for tool composition: types carry convention bundles (per-dimension frames); a declared tool posts a known transport between frames while an opaque tool emits a value of latent (undeclared) convention; reconciliation posts the constraint that two frames agree; and disclosure declares a latent frame. The typing context, read as a presheaf over the composition poset, is a semantic interface complex. Coherence decomposes into two conditions that must not be conflated: consistency (the declared constraints and reconciliations have a common solution — no clash) and  $\text{fee} = 0$  (every reconciliation is forced by the declarations — no undeclared convention is load-bearing), where  $\text{fee} = \text{rank}(0 \cup R) - \text{rank}(0)$  for  $0$  the declared constraints and  $R$  the reconciliations —

the companion’s witness rank, computed as the deployed checker computes it. A coercion is an admissible disclosure; the minimum coercion set is the disclosure normal form, of size fee. Part II (dynamics and soundness) gives a tagged operational semantics in which a reconcile whose two runtime frames disagree is convention-stuck, and proves Coherence Soundness: a coherent program enjoys progress — it is a value or steps, and never reaches a convention-stuck state, for every behavior of the opaque tools — and preservation — reduction preserves typing and coherence. The progress proof is the payoff and it is short once the statics are right: at runtime the declared constraints  $\emptyset$  hold by construction, and  $\text{fee} = \emptyset$  makes every reconciliation a linear consequence of  $\emptyset$ , so consistency forces it to hold. The calculus is Hindley–Milner unification with a cohomological twist — a cycle of frame constraints is satisfiable iff its holonomy is trivial — and its soundness theorem is the positive half of the program’s thesis: coherence is a global typing invariant precisely because it is locally unlearnable.

---

## 1 1. Introduction

### 1.1 1.1 The missing composition rule, as a typing problem

Hoare’s sequential-composition rule

$$\frac{\{P\} C_1 \{R\} \quad \{R\} C_2 \{Q\}}{\{P\} C_1; C_2 \{Q\}}$$

is sound because the assertion  $R$  means the same thing where  $C_1$  leaves off and where  $C_2$  picks up. Open compositions violate this silently: a calendar tool emits a date in ISO-8601 with implicit UTC; an invoicing tool reads it in implicit local time; each interface is type-correct, each local contract holds, yet the composition is wrong by a timezone offset, invisible to either tool’s local audit ([CD] §1; provably invisible to any bounded-locality observation, [PAR]). The seam moved the semantic frame — the convention under which a value is read.

The companion program names the obstruction. A composition is a semantic interface complex ([CD] Def 2.1): a presheaf  $F$  of conventions over a poset of contexts, an observable subpresheaf  $\emptyset \subseteq F$  (declared conventions), and a latent quotient  $L = F/\emptyset$  (undeclared ones). The coherence fee is the witness rank, [CD]’s unique invariant respecting the primitive of repair (disclosure). This paper makes the apparatus a type system: the missing composition rule is the missing typing rule — the side condition that a composition’s conventions glue.

### 1.2 1.2 The thesis, stated precisely

We build  $\lambda_\nabla$  (“lambda-connection”). A term generates, per dimension  $d$ , a system of group-valued constraints on the frames (conventions) of its value occurrences:

- a declared tool posts a known relation  $x_{\text{out}} = g + x_{\text{in}}$  (its declared transport);

- an opaque tool emits a value whose frame is latent — a fresh free variable, posting no relation (the tool does not declare its convention);
- a reconciliation  $\text{reconcile}_d(v_1, v_2)$  posts the constraint  $x_{\{v_1\}} = x_{\{v_2\}}$  — the seam, the only former (in the recursion-free core) that demands two independently-derived frames agree;
- a disclosure  $\text{disclose}_d(v, \varphi)$  posts a pin  $x_v = \varphi$ , declaring a latent frame.

Write  $\mathcal{O}$  for the declared constraints (declared transports and pins) and  $\mathcal{R}$  for the reconciliation constraints. The central definitions:

Coherence  $\Gamma \vdash e \checkmark : \Leftrightarrow \mathcal{O} \cup \mathcal{R}$  is consistent (no clash) and  $\text{fee}(e) = \emptyset$ , where  $\text{fee}(e) := \text{rank}(\mathcal{O} \cup \mathcal{R}) - \text{rank}(\mathcal{O})$  is the number of reconciliations not forced by the declarations.

Two failure modes, genuinely different and not to be conflated:

- Inconsistency (a clash):  $\mathcal{O} \cup \mathcal{R}$  has no solution — the declared transports and a reconciliation contradict (a settlement loop whose declared arithmetic does not compose; or two pins reconciled at mismatched frames). A hard type error;  $\text{fee}$  (a dimension) does not see it.
- Under-determination (positive fee):  $\mathcal{O} \cup \mathcal{R}$  is solvable but some reconciliation is not a consequence of  $\mathcal{O}$  — it constrains a latent frame the declarations leave free, so at run time that frame (set by an opaque tool) can violate it. Repairable by disclosure.

The dictionary, each entry made precise in §2–§3:

Type-theoretic notion	Companion object
A type's convention bundle	A stalk of the convention presheaf $F$ ([CD] Def 2.1)
Typing context $\Gamma$ , as a presheaf over the composition poset	A semantic interface complex $G_e \in \text{SemComp}$
Declared transport / opaque tool	Observable restriction map / a latent stalk in $L = F/\mathcal{O}$
Coherence	Consistency and $\text{fee} = \emptyset$
Convention clash (hard error)	Inconsistent declared constraints ( $\mathcal{O} \cup \mathcal{R}$ unsatisfiable)
Coherence fee	Witness rank $\text{rank}(\mathcal{O} \cup \mathcal{R}) - \text{rank}(\mathcal{O}) = \dim H^1$ ([CD])
A coercion	An admissible disclosure ([CD] §2.3 III)
Minimum coercion set	The disclosure normal form ([CD] Thm 7.1), size $\text{fee}$
Convention-mismatch stuck state (Part II)	Holonomy $\neq \emptyset$ realized at run time

The one-line summary:  $\lambda_{\nabla}$  is Hindley–Milner unification with a cohomological twist. In ordinary HM a cycle of equality constraints  $\alpha = \beta = \dots = \alpha$  is always satisfiable. In  $\lambda_{\nabla}$  the constraints carry group elements, and a cycle is satisfiable iff its holonomy (the signed sum of transports) is  $\emptyset$ ; a reconciliation that closes such a cycle is forced only when the declarations fix that holonomy. This is the exact sense in which the separation-logic frame rule fails ([CD] §6.3 calls the fee “the minimum disclosure required to recover frameability”).

### 1.3 1.3 Why a type system: the bridge from the hardness result

[PAR] proves the obstruction is a parity, hence (i) invisible to bounded-locality observation and (ii) hard to recover from noisy observation under LPN/LWE. The constructive contrapositive is this paper’s reason to exist:

Coherence cannot be recovered bottom-up from observation; therefore it must be carried top-down, as a typing obligation declared at the interface.

$\lambda_{\nabla}$  is that declared structure; the certificate is the typing derivation. B2 (negative) and B3 (positive) are a matched pair — coherence is a global typing invariant precisely because it is locally unlearnable.

### 1.4 1.4 Contributions

Part I (§2–§3). The calculus  $\lambda_{\nabla}$  (recursion-free core; §2); its convention system  $(0, R)$ ; the constraint-generation statics (§3.1); the consistency / fee decomposition and the coherence judgment (§3.2–3.3); decidability of coherence as a polynomial-time rank (§3.2); the exact-regime conditions making the fee the companion’s witness rank (§3.3); coercion = admissible disclosure and minimum coercion = disclosure normal form (§3.4); frame polymorphism (§3.5).

Part II (§4–§5). A tagged operational semantics with the convention-stuck state (§4); the main result, Coherence Soundness — progress (coherent  $\Rightarrow$  value-or-steps, never convention-stuck, for every opaque behavior) and preservation (§5) — with the elaboration-soundness corollary (elaboration inserts fee coercions and yields a sound program).

Deferred: the graded disclosure modality  $\square_r$  (grade = fee), general feedback/recursion, and interface-evolution-as-preservation ([CD] Conj 9.5) are Part III; mechanization and the Bulla elaborator are Part IV.

### 1.5 1.5 Claim discipline and scope

[Established, cited]: the semantic-interface-complex model, the witness rank and its TU incidence-rank computation (Bulla), repair duality (Thm 6.2), the disclosure normal form (Thm 7.1), the DFD+CHP  $\Rightarrow$  exact sufficiency — from [CD] and the hierarchical-fee companion [HF]; the holonomy / flatness classification from [LV]; the hardness motivation from [PAR]. [New,

proven here]:  $\lambda_{\nabla}$  and its statics; the consistency/fee decomposition and coherence characterization (Prop 3.5); the exact-regime conditions for  $\lambda_{\nabla}$ 's complexes (Prop 3.3); coercion = disclosure (Prop 3.6); decidability (Prop 3.2); and the Coherence Soundness theorem (Thm 5.4: progress + preservation). Scope: the abelian, exact, recursion-free regime — each dimension has an abelian additive structure group; constraints are single-dimension and rank-1; declared transports are fixed (not frame-dependent); the core has no recursion (reconciliation is the sole loop-closer). Frame-dependent transports, cross-dimension coupling, multi-frame interfaces (obstruction in  $\mathbb{H}^2$ /gerbes, [SCPI]), and general feedback are the frontier, named and out of scope (§3.7).

---

## 2 2. The calculus $\lambda_{\nabla}$

### 2.1 2.1 Frames, declared transports, and opacity

Fix a finite set  $D$  of convention dimensions. For each  $d$  fix an abelian group  $\mathbb{E}_d$ , written additively, the structure group, acting freely and transitively on a nonempty set  $\mathbb{F}_d$  of frames: two frames differ by a unique transport  $\psi - \phi \in \mathbb{E}_d$ . (Multiplicative groups such as unit-scaling are presented additively via a fixed logarithm; we restrict to groups admitting such a presentation, so the incidence computation of §3.2 is over  $\mathbb{Q}$  and totally unimodular.)

A tool transforms a value. On each relevant dimension  $d$  it is either:

- declared: it posts a known transport  $g^{\wedge}t_d \in \mathbb{E}_d$ , so its output frame is  $g^{\wedge}t_d +$  its input frame (a known relation between the two frames); or
- opaque: it does not declare its convention behavior, so its output frame is a latent frame — a fresh free variable, related to nothing. The calendar tool emitting “implicit UTC” is opaque on `timezone`: it produces a date whose frame is undeclared.

Latent frames are the hidden data the fee measures; a latent frame becomes observable only by disclosure. This is [CD]'s  $0 \subseteq \mathbb{F}$  split: declared frames generate the observable  $0$ , latent ones the quotient  $\mathbb{L} = \mathbb{F}/0$ . Note transports of declared tools are known and fixed — coherence is not about discovering them but about whether the declared structure forces the reconciliations.

### 2.2 2.2 Types

Types carry, on relevant dimensions, a frame term  $\phi$  — a concrete frame or a frame variable  $\alpha$ :

$$\tau ::= b\langle d : \phi \rangle \mid \tau_1 \rightarrow \tau_2 \mid \tau_1 \times \tau_2 \mid \text{unit}, \quad \varphi ::= \phi \in \mathbb{F}_d \mid \alpha.$$

$b\langle d : \phi \rangle$  is a single-dimension interface type (one relevant  $d$  per base value; a value relevant to several dimensions is a product). The single-dimension restriction makes the per-dimension constraint systems independent (Prop 3.3); multi-frame interfaces, where dimensions couple,

are the non-abelian frontier (§3.7). A frame variable is observable if a derivation pins it, latent otherwise.

### 2.3 2.3 Terms

$\lambda_\nabla$  is the simply-typed  $\lambda$ -calculus — without recursion or fixpoint (§3.7) — plus tools, a binder, reconciliation, and disclosure:

$$\begin{aligned}
e ::= & x \mid () \mid (e_1, e_2) \mid \pi_i e \mid \lambda x:\tau. e \mid e_1 e_2 \mid \text{let } x=e_1 \text{ in } e_2 \\
& \mid t(e) \quad (\text{tool: declared posts a transport, opaque posts a latent frame}) \\
& \mid \text{reconcile}_d(e_1, e_2) \quad (\text{seam: assert the two frames agree}) \\
& \mid \text{disclose}_d(e, \phi) \quad (\text{coercion: declare } e\text{'s } d\text{-frame is } \phi).
\end{aligned}$$

$\text{reconcile}_d(e_1, e_2)$  returns the common value and posts  $x_{\{e_1\}} = x_{\{e_2\}}$  — the only former of the recursion-free core that can demand two independently-built frames agree (Prop 2.2).  $\text{disclose}_d(e, \phi)$  pins  $e$ 's frame to  $\phi$ ; in elaboration (§3.4) disclosures are inserted automatically.

### 2.4 2.4 The convention system of a term; the context as presheaf

A well-scoped term  $e$  generates, per dimension  $d$ , a finite set of frame variables (one per  $d$ -relevant value occurrence) and two sets of group-valued affine constraints over them:

- the declared constraints  $0_{\{e, d\}}$ : one  $x_v = g^{\wedge t_d} + x_u$  per declared tool edge  $u \rightarrow v$ , and one pin  $x_v = \phi$  per disclosure;
- the reconciliation constraints  $R_{\{e, d\}}$ : one  $x_{\{v_1\}} = x_{\{v_2\}}$  per  $\text{reconcile}_d$ .

Opaque tools post no constraint (their output is a fresh free variable). Collecting these over  $d$  gives the convention sheaf  $G_e \in \text{SemComp}$  ([CD] Def 2.1): contexts are value occurrences ordered by dataflow; the stalk at a node is its frame; declared edges are restriction maps; the observable subpresheaf  $0_e$  is generated by the declared constraints; the latent  $L_e = F_e/0_e$  is carried by the opaque outputs. The typing context  $\Gamma$ , read as a presheaf over the composition poset, is  $G_e$ .

Definition 2.1 (holonomy). For a cycle in the graph of declared + reconciliation edges, its holonomy is the signed sum of transports along it (reconciliation edges contribute  $\emptyset$ ). A cycle is coherent iff its holonomy is  $\emptyset$ . (This is [LV] Def 2.3, additively.)

Proposition 2.2 (reconciliation is the sole loop-closer). In the recursion-free core, the dataflow is a DAG, so the declared/opaque tool edges form no cycle; every cycle in  $G_e$  uses a reconciliation edge or two pins meeting at a pinned frame.

Proof. Without `fix`, `let`/application build a finite DAG of value dependencies; tool edges follow it and form no cycle. The only edges against the DAG are reconciliation edges and pins. ■

## 2.5 Two running examples

(a) Settlement — a clash. The pipeline of [LV] Ex 2.9, with risk and treasury declared on the settlement–instant dimension (known day-arithmetic transports  $g_r, g_t$ ):

```
let cap = capture()           -- cap : Date(d:α)           (α a frame variable)
let h   = risk(cap)          -- declared: posts  x_h = g_r + x_cap
let rel = treasury(h)        -- declared: posts  x_rel = g_t + x_h
in reconcile_d(cap, rel)     -- R: posts  x_cap = x_rel
```

$0$  forces  $x_{rel} - x_{cap} = g_t + g_r$ ;  $R$  demands  $x_{rel} - x_{cap} = 0$ . So  $R$  is a consequence of  $0$  ( $fee = 0$ ), but  $0 \cup R$  is inconsistent unless  $g_t + g_r = 0$ . The day-arithmetic is off by one banking day, so  $g_t + g_r \neq 0$ : a clash, a hard type error — the declared transports genuinely do not compose. (No disclosure repairs a clash; one must fix the arithmetic. [LV] §2.9 instead treats the basis as frame-dependent — the transport depends on the undeclared basis — which is the cross-dependency frontier of §3.7; there the same situation is  $fee = 1$ . In the scoped fixed-transport model it is a clash, and we report it as one.)

(b) An undeclared convention —  $fee = 1$ . A genuinely under-determined reconciliation:

```
let a = fetch_rate()         -- declared: a : Rate(d:φ_usd) (pinned, observable)
let b = legacy_quote()       -- OPAQUE on d: b : Rate(d:β)   (β latent -
legacy tool, undeclared scale)
in reconcile_d(a, b)         -- R: posts  x_a = x_b
```

$0 = \{x_a = \phi\_usd\}$ ;  $R = \{x_a = x_b\}$  with  $x_b$  a latent free variable (opaque output).  $0 \cup R$  is consistent (set  $x_b = \phi\_usd$ ), but  $R$  is not a consequence of  $0$  (it constrains the free  $x_b$ ):  $fee = rank(0 \cup R) - rank(0) = 2 - 1 = 1$ . The legacy quote’s scale is undeclared, so the type-checker cannot certify the rates agree. One coercion repairs it — `disclose_d(b, φ_usd)` declares the scale, pinning  $x_b$ ; then  $R$  follows from  $0$ ,  $fee = 0$ , and the term is coherent. One obstruction, one coercion:  $fee = \rho = 1$ , repair duality ([CD] Thm 6.2).

---

## 3 3. Statics

### 3.1 3.1 Constraint generation

$\Gamma \vdash e : \tau \rightsquigarrow (0; R)$  emits the declared ( $0$ ) and reconciliation ( $R$ ) constraints, under a freshness discipline (each value occurrence draws frame variables from a global counter, so  $(0; R)$  is canonical). The structural rules (Var, Unit, Pair, Proj, Abs, App, Let) thread frame variables unchanged and emit nothing; App matches the argument’s frame terms to the parameter’s, emitting only forward identity constraints into  $0$  (no back-edge: by Prop 2.2 the argument is built earlier in the DAG). Let generalizes (frame variables and transports) via `Gen_Γ` (§3.5). The characteristic rules:

$$\frac{\Gamma \vdash e : \mathbf{b}\langle d:\alpha \rangle \rightsquigarrow (O; R) \quad t \text{ declared, } g_d^t}{\Gamma \vdash t(e) : \mathbf{b}'\langle d:\beta \rangle \rightsquigarrow (O \cup \{\beta = g_d^t + \alpha\}; R), \beta \text{ fresh}} \text{ (Tool-decl)}$$

$$\frac{\Gamma \vdash e : \mathbf{b}\langle d:\alpha \rangle \rightsquigarrow (O; R) \quad t \text{ opaque on } d}{\Gamma \vdash t(e) : \mathbf{b}'\langle d:\beta \rangle \rightsquigarrow (O; R), \beta \text{ fresh (latent — no constraint)}} \text{ (Tool-opaque)}$$

$$\frac{\Gamma \vdash e_1 : \mathbf{b}\langle d:\alpha \rangle \rightsquigarrow (O_1; R_1) \quad \Gamma \vdash e_2 : \mathbf{b}\langle d:\beta \rangle \rightsquigarrow (O_2; R_2)}{\Gamma \vdash \text{reconcile}_d(e_1, e_2) : \mathbf{b}\langle d:\alpha \rangle \rightsquigarrow (O_1 \cup O_2; R_1 \cup R_2 \cup \{\alpha = \beta\})} \text{ (Recon)}$$

$$\frac{\Gamma \vdash e : \mathbf{b}\langle d:\alpha \rangle \rightsquigarrow (O; R)}{\Gamma \vdash \text{disclose}_d(e, \phi) : \mathbf{b}\langle d:\phi \rangle \rightsquigarrow (O \cup \{\alpha = \phi\}; R)} \text{ (Disclose)}$$

(Tool-decl) adds a declared relation to 0; (Tool-opaque) adds nothing (latent output); (Recon) adds the seam constraint to R; (Disclose) adds a pin to 0. Generation is syntax-directed, linear in  $|e|$ .

### 3.2 3.2 Consistency, the fee, and decidability

Treat  $O$  and  $R$  as affine systems  $M_O \times = \mathbf{b}_O, M_R \times = \mathbf{b}_R$  over  $\mathbb{Q}^{\{D\}}$  (one block per dimension; the rows of  $M_O$  are declared transports and pins, of  $M_R$  the reconciliations; entries are  $\{-1, 0, +1\}$  signed-incidence scaled by transports, with offsets  $\mathbf{b}$ ).

Definition 3.1 (consistency; fee).  $e$  is consistent iff the combined system  $O \cup R$  has a solution. The coherence fee is

$$\text{fee}(e) := \text{rank} \begin{bmatrix} M_O \\ M_R \end{bmatrix} - \text{rank}(M_O),$$

the number of reconciliation constraints linearly independent of the declared constraints — equivalently  $\dim(R \bmod \text{span } O)$ , the reconciliations not forced by the declarations. This is [CD]'s witness rank  $\text{rank}(\delta_F) - \text{rank}(\delta_O)$  under the identification  $O = \delta_O$  (observable),  $O \cup R$ -row-space =  $\delta_F$  (full). It is the witness rank the deployed checker reports ([CD] §10.1); that checker computes it via the Kron-reduced Gram, which equals this incidence-rank difference in the exact regime ([CD] §3.5).  $\lambda_\nabla$  uses only the elementary incidence form (§3.3).

Consistency is a value condition (do the declarations and reconciliations have a common solution?); the fee is a structural dimension (how many reconciliations are not forced?). They are independent: a clash can occur at  $\text{fee} = 0$  (settlement:  $R \subseteq \text{span } O$  but the RHS contradicts — §2.5a), and  $\text{fee} > 0$  can occur with no clash (an undeclared convention, satisfiable but not forced — §2.5b). A term whose every tool is declared on every relevant dimension has  $\text{fee} = 0$  (no latent frames), so the fee machinery (§3.4) fires exactly on the opaque tools that model undeclared conventions.

Proposition 3.2 (decidability). Consistency and  $\text{fee}(e)$  are decidable in time  $O(|e|^\omega)$  over  $\mathbb{Q}$ .

Proof. Both are rank computations on the  $\{-1, 0, +1\}$ -incidence-with-transport matrices  $M_0, [M_0; M_R]$ :  $\text{fee}$  is the rank difference; consistency is Rouché–Capelli ( $\text{rank}[M_{\{0 \cup R\}}] = \text{rank}$  of its augmented matrix  $[M_{\{0 \cup R\}} \setminus, | \setminus, b]$ ). Gaussian elimination over  $\mathbb{Q}$  is polynomial; the unweighted incidence pattern is totally unimodular, so the  $\text{fee}$  rank is field-independent (matching [CD]’s signed-incidence TU, [SI]). ■

Proposition 3.3 ( $\lambda_\nabla$  complexes are exact; the  $\text{fee}$  is the witness rank). Under the core’s discipline — single-dimension interface types, one frame per node, per-dimension constraints — the per-dimension systems are mutually independent and rank-1, so  $G_e$  satisfies disjoint field decomposition and the class-homogeneous partition of [HF]; hence  $G_e$  lies in [CD]’s exact regime and  $\text{fee}(e)$  of Def 3.1 equals the witness rank  $\dim H^1_{\text{seam}}(G_e)$ .

Proof. Every rule of §3.1 names a single dimension and the single-dimension interface forbids a node bearing two frames, so the constraint set partitions across dimensions with no coupling (DFD), each node carrying one frame (CHP). By [HF] (Theorem 2: DFD+CHP  $\Rightarrow$  exact regime — proved there, not in [CD]), the seam complex truncates to two terms and the witness rank equals the two-term incidence-rank formula, which is exactly  $\text{rank}[M_0; M_R] - \text{rank } M_0$  ([CD] §3.5, the Bulla computation). ■

We make no claim about [CD]’s Kron/Schur identity (Lemma 3.9);  $\lambda_\nabla$  uses the elementary incidence-rank form, which is all the checker needs.

### 3.3 The coherence judgment

Definition 3.4 (coherence).  $\Gamma \vdash e : \tau \checkmark : \iff \Gamma \vdash e : \tau \rightsquigarrow (0; R), 0 \cup R$  is consistent, and  $\text{fee}(e) = 0$ .

Proposition 3.5 (coherence = the declarations force every reconciliation).  $\Gamma \vdash e : \tau \checkmark$  iff every reconciliation constraint in  $R$  is a linear consequence of the declared constraints  $0$  with matching right-hand side — i.e.  $R \subseteq \text{span}(0)$  ( $\text{fee} = 0$ ) and the forced value of each reconciliation is  $0$  (consistency). Equivalently: the declared transports and pins determine every reconciliation and force its holonomy to  $0$ .

Proof.  $\text{fee} = 0$  says  $\text{rank}[M_0; M_R] = \text{rank } M_0$ , i.e. each row of  $M_R$  lies in the row space of  $M_0$  — every reconciliation’s coefficient vector is a combination  $\sum c_i (M_0)_i$ . The scalar  $\sum c_i (b_0)_i$  is well-defined (independent of the non-unique choice of  $c_i$  when the  $M_0$  rows are dependent): whenever  $0$  is satisfiable, fixing any solution  $\xi$  of  $M_0 \xi = b_0$  gives  $\sum c_i (b_0)_i = (\sum c_i (M_0)_i) \setminus, \xi = M_r \setminus, \xi$ , a value determined by  $M_r$  alone — this is the forced holonomy of the reconciliation, well-posed exactly on the  $0$ -consistent branch. Consistency of  $0 \cup R$  then forces it to agree with the reconciliation’s RHS:  $M_r \xi' = 0$  for a witness  $\xi'$  of  $0 \cup R$ , and since  $\xi'$  also satisfies  $0$ ,  $\sum c_i (b_0)_i = M_r \xi' = 0$ . Conversely if either fails,  $R \not\subseteq \text{span } 0$  ( $\text{fee} > 0$ ) or some forced holonomy is  $\neq 0$  (clash). ■

So coherence is precisely “the declared conventions determine every reconciliation and force it to close.” A clash breaks force-it-to-close (a determined holonomy is nonzero); a positive fee breaks determine (a reconciliation runs through an undeclared convention) — and only the latter is repairable by declaring more (§3.4). This is the condition Part II’s progress theorem will discharge.

### 3.4 3.4 Coercion is disclosure; minimum coercion is the fee

A disclosure declares a line in the observable presheaf — and the line may be relational or absolute. The absolute pin  $\text{disclose}_d(e, \varphi)$  posts  $x_e = \varphi$ ; the relational disclosure  $\text{disclose}_d(e_1, e_2, g)$  posts  $x_{\{e_1\}} - x_{\{e_2\}} = g$ , declaring the convention relationship between two values (e.g. “these two rates are in the same scale”,  $g = 0$ ). Both are admissible disclosures ([CD] §2.3 III) when consistent with the frames already forced below them (downward-closure), which the checker verifies; [CD] Lemma 5.3 guarantees an admissible representative at any minimal latent context. The relational form is essential: the obstruction classes are reconciliation differences  $x_{\{e_1\}} - x_{\{e_2\}}$ , so the canonical coercion discloses a class (a line in  $H^1$ ), which is generally relational. Absolute pins alone do not suffice for repair duality — pinning columns can require more than fee disclosures when a reconciliation relates two free frames — whereas disclosing the fee obstruction classes always reaches coherence in exactly fee steps (Thm 3.7).

Proposition 3.6 (coercion = admissible disclosure). Each downward-closed disclosure (absolute or relational) induces an admissible one-generator disclosure of  $G_e$ ; minimal-context disclosures always have an admissible representative ([CD] Lemma 5.3). ■

Elaboration repairs under-determination, not clashes. When  $e$  is consistent but  $\text{fee}(e) > 0$ , the checker elaborates  $\Gamma \vdash e \rightsquigarrow e' : \tau$  by inserting coercions until  $\text{fee} = 0$ . (A clash is not repaired by disclosure — declaring more frames cannot remove a contradiction; it is reported as a hard error.)

Theorem 3.7 (minimum coercion = disclosure normal form). For consistent  $e$ , a coercion set making  $e$  coherent exists; the minimum has size exactly  $\text{fee}(e)$ , and a minimum set is a basis of the witness matroid — the disclosure normal form; a minimum-cost set under a per-field cost is the min-weight basis, by the matroid greedy in  $O(\text{fee}(e) \cdot |e|^2)$ .

Proof.  $G_e \in \text{SemComp}_{\text{ex}}$  (Prop 3.3) and  $e$  consistent, so [CD] Thm 6.2 (repair duality  $\rho = r$ ) and Thm 7.1 (disclosure normal form) apply via Prop 3.6; each admissible nontrivial disclosure drops fee by one ([CD] Lemma 3.7) and (being consistent with the forced frames) introduces no clash; the matroid-greedy minimum-cost statement is [LV]’s repair theorem. ■

### 3.5 3.5 Frame polymorphism

$\text{Gen}_\Gamma$  generalizes frame variables and transports not free in  $\Gamma$ . The pipe combinator is typed once:

$$\text{pipe} : \forall \alpha \forall g \forall g'. (\mathbf{b}\langle d:\alpha \rangle \rightarrow \mathbf{b}'\langle d:g+\alpha \rangle) \rightarrow (\mathbf{b}'\langle d:g+\alpha \rangle \rightarrow \mathbf{b}''\langle d:g'+g+\alpha \rangle) \rightarrow \mathbf{b}\langle d:\alpha \rangle \rightarrow \mathbf{b}''\langle d:g'+g+\alpha \rangle,$$

$\text{pipe } f \ g \ x = g \ (f \ x)$  (transports  $\forall$ -bound, determined at the call site).  $\text{pipe}$  posts no reconciliation, so by Prop 2.2 it generates no obstruction and instantiates soundly at any frames.

Proposition 3.8 (polymorphism hides no obstruction, in the recursion-free core). A closed, reconciliation-free combinator has  $\text{fee} = \emptyset$  and is consistent at every instantiation; obstruction arises only when its result is later fed into a  $\text{reconcile}$ . Hence generalization is sound.

Proof. By Prop 2.2 a reconciliation-free term has acyclic  $G_e$  and empty  $R$ , so  $\text{fee} = \text{rank}[M_0] - \text{rank } M_0 = \emptyset$  and  $\emptyset$  (a triangular DAG system) is consistent. A later  $\text{reconcile}$  adds the constraint at its own use site. ■

(Obstruction is a property of loop-closing use sites, not of polymorphic combinators. A combinator that does close a loop — a  $\text{feedback}/\text{fix}$  — carries a residual coherence constraint in its scheme; that is the graded-modality content of Part III, and why the core excludes it.)

### 3.6 3.6 The frame-rule reading

Separation logic's frame rule is sound when the framed assertion is stable;  $\lambda_\nabla$ 's composition is sound when the conventions glue — coherent. [CD] §6.3 names the  $\text{fee}$  “the minimum disclosure required to recover frameability”; Thm 3.7 makes it operational.

### 3.7 3.7 The boundaries

The following are all out of scope (§1.5): (i) recursion/feedback (a  $\text{fix}$  closes a loop through application, breaking Prop 2.2 — Part III's graded modality); (ii) frame-dependent transports (a declared transport depending on another dimension's frame — the settlement basis, §2.5a — couples dimensions); (iii) multi-frame interfaces / non-thin coercions, where the obstruction is in  $\mathbb{H}^2$  (a gerbe, [SCPI]) and coercions must themselves cohere — a 2-categorical typing problem.  $\lambda_\nabla$  is the 1-categorical, abelian, recursion-free core.

## 4 Part II. Dynamics and Soundness

### 5 4. Operational semantics

We give a call-by-value small-step semantics on frame-tagged values. The point of Part II is that the static coherence judgment (§3.3) genuinely guarantees the dynamic property it was designed for: a coherent program never gets stuck reconciling mismatched conventions, no matter what the opaque tools actually do.

#### 5.1 4.1 Runtime configurations

Runtime values carry concrete frames:

$$v ::= \langle k \rangle_{d:\phi} \mid (v_1, v_2) \mid \langle \lambda x:\tau. e, \rho \rangle,$$

where  $\langle k \rangle_{d:\phi}$  is a base datum  $k$  tagged with a concrete frame  $\phi \in \mathbb{F}_d$  on its relevant dimension, and  $\rho$  is a value environment. The opaque tools' actual behavior is fixed by an opacity environment  $\varepsilon$ , assigning to each opaque-tool occurrence  $o$  and dimension  $d$  a concrete actual transport  $\varepsilon(o, d) \in \mathbb{E}_d$  — what the tool really does, which the program never declared. Reduction  $\rightarrow_\varepsilon$  is parameterized by  $\varepsilon$ ; the soundness theorem quantifies over all  $\varepsilon$ .

#### 5.2 4.2 Reduction

The standard rules ( $\beta$ , let, projection) are as usual. The convention-relevant rules:

- Declared tool.  $t(\langle k \rangle_{d:\phi}) \rightarrow_\varepsilon \langle k \rangle_{d:\phi} \circ g^{\wedge}_d$  — the output frame is the input frame transported by the declared  $g^{\wedge}_d$ .
- Opaque tool at occurrence  $o$ .  $t(\langle k \rangle_{d:\phi}) \rightarrow_\varepsilon \langle k \rangle_{d:\phi} \circ \varepsilon(o, d)$  — the output frame is shifted by the tool's actual transport  $\varepsilon(o, d)$ , which the type system did not know. (Equivalently, the opaque output's frame is an arbitrary  $\mathbb{E}_d$ -translate;  $\varepsilon$  ranges over all.)
- Disclosure.  $\text{disclose}_d(\langle k \rangle_{d:\phi}, \psi) \rightarrow_\varepsilon \langle k \rangle_{d:\psi}$  — a trusted re-tag to the declared frame  $\psi$ . (Disclosure is the programmer's assertion of a convention; the type system's job is to make reconciliations succeed given the declarations. A mis-declaration is garbage-in, outside the theorem.)
- Reconciliation.

$$\text{reconcile}_d(\langle k_1 \rangle_{d:\phi_1}, \langle k_2 \rangle_{d:\phi_2}) \rightarrow_\varepsilon \begin{cases} \langle k_1 \rangle_{d:\phi_1} & \text{if } \phi_1 = \phi_2, \\ \text{stuck}_d & \text{if } \phi_1 \neq \phi_2. \end{cases}$$

**Definition 4.1 (convention-stuck).** A closed term is convention-stuck if it is not a value and its only redex is a  $\text{reconcile}_d(v_1, v_2)$  with  $v_1, v_2$  carrying unequal  $d$ -frames. (Every other

non-value closed well-typed term has a redex, by the standard progress argument for STLC; the convention-stuck state is the only new stuck form  $\lambda_{\nabla}$  introduces.)

## 6 5. Soundness

The whole-system frame state at any point of a run is captured by a valuation: assign to each frame variable  $x_v$  (a value occurrence) the concrete frame its runtime value carries. We show this valuation always satisfies the declared constraints, and that coherence then forces the reconciliations.

Lemma 5.1 (the declared constraints hold at run time). Let  $e$  be well-typed and  $\eta$  the valuation induced by any reduction state of  $e$  under any  $\varepsilon$ . Then  $\eta$  satisfies every declared constraint in  $0$  (every declared-tool relation and every pin).

Proof. By cases on the reduction rules. A declared tool produces  $\langle \cdot \rangle_{d: g^t_d + \phi}$  from input  $\langle \cdot \rangle_{d: \phi}$ , so  $\eta(x_{out}) = g^t_d + \eta(x_{in})$  — the declared relation  $x_{out} = g^t_d + x_{in}$  holds. A disclosure re-tags to  $\psi$ , so  $\eta(x_v) = \psi$  — the pin  $x_v = \psi$  holds. Opaque tools post no  $0$ -constraint, so  $\varepsilon$ 's arbitrary choices constrain nothing in  $0$ . Substitution ( $\beta$ ,  $\text{let}$ ) carries a value with its frame intact, preserving the relations. Hence  $\eta \models 0$  throughout. ■

Note  $\eta$  need not satisfy  $R$ : an opaque tool's actual transport  $\varepsilon(o, d)$  can make a reconciliation's two frames differ. That is exactly the convention-stuck risk; coherence rules it out.

Lemma 5.2 (coherence forces the reconciliations). If  $\vdash e \checkmark$  (consistent,  $\text{fee} = 0$ ) then every valuation  $\eta$  with  $\eta \models 0$  also satisfies every reconciliation in  $R$ .

Proof. By Prop 3.5,  $\text{fee} = 0$  gives  $R \subseteq \text{span}(0)$ : each reconciliation  $r \in R$  is  $M_r = \sum_i c_i (M_0)_i$  for scalars  $c_i$ . For any  $\eta \models 0$  we have  $(M_0)_i \setminus \eta = (b_0)_i$ , so  $M_r \setminus \eta = \sum_i c_i (b_0)_i$ . Consistency (Prop 3.5) gives  $\sum_i c_i (b_0)_i = b_r = 0$ . Hence  $M_r \setminus \eta = 0$ , i.e.  $\eta$  satisfies  $r$ . ■

Lemma 5.3 (preservation). If  $\vdash e : \tau \checkmark$  and  $e \rightarrow_{\varepsilon} e'$  then  $\vdash e' : \tau \checkmark$  (for every  $\varepsilon$ ).

Proof. Typing preservation is the standard STLC substitution argument, extended to the tool/disclose/reconcile forms; we check coherence is preserved. Each step either (i) realizes a declared relation (a declared tool fires: the relation  $x_{out} = g + x_{in}$  becomes a fact about concrete frames, which only adds to  $0$ 's consequences —  $R \subseteq \text{span } 0$  is preserved and no clash is created since by Lemma 5.1 the realized value satisfies  $0$ ); (ii) re-tags by a disclosure (likewise adds a pin already in  $0$ ); (iii) discharges a reconciliation that, by Lemma 5.2, held — removing a satisfied  $R$ -constraint cannot increase  $\text{fee}$  or create inconsistency; or (iv) is a  $\beta/\text{let}/\text{projection}$  step, whose substitution lemma (a value of frame  $\phi$  substituted for a variable of frame term  $\alpha$  adds the pin  $\alpha = \phi$ , consistent with  $0$  by Lemma 5.1) preserves both  $R \subseteq \text{span } 0$

and consistency. An opaque tool firing adds no constraint, so trivially preserves coherence. In every case  $\text{fee}(e') = 0$  and  $e'$  consistent. ■

Theorem 5.4 (Coherence Soundness). Let  $\vdash e : \tau \checkmark$  be a closed coherent term. Then for every opacity environment  $\varepsilon$ :

1. (Progress)  $e$  is a value or  $e \rightarrow_{\varepsilon} e'$  for some  $e'$ ; and  $e$  is never convention-stuck.
2. (Preservation) if  $e \rightarrow_{\varepsilon} e'$  then  $\vdash e' : \tau \checkmark$ .

Hence no reduction sequence of a coherent term reaches `stuck_d` — whatever the opaque tools do.

Proof. (2) is Lemma 5.3. For (1): by the standard STLC progress argument, a closed well-typed non-value term has a redex; the only redex that could fail to step is a `reconcile_d(v_1, v_2)` with  $v_1, v_2$  carrying unequal frames (Def 4.1). Suppose such a redex is reached. By preservation (2) the reached term is still coherent, so its reconciliation constraint  $r$  is in  $R$ . By Lemma 5.1 the runtime valuation  $\eta$  satisfies  $0$ ; by Lemma 5.2 it then satisfies  $r$ , i.e.  $\eta(x_{\{v_1\}}) = \eta(x_{\{v_2\}})$  — the two frames are equal, contradicting the supposed mismatch. So no convention-stuck redex is ever reached, and the term steps. ■

The shape of the proof is the point. Progress is short because the statics are right: the declarations hold at run time unconditionally (Lemma 5.1), and  $\text{fee} = 0$  makes every reconciliation a linear consequence of the declarations (Lemma 5.2), so consistency forces it. The opaque tools' arbitrary behavior ( $\varepsilon$ ) enters only off the forced reconciliations — exactly the content of  $\text{fee} = 0$ . This is the operational meaning of the companion's slogan that the fee is "the minimum disclosure required to recover frameability": with  $\text{fee} = 0$  the frame is recovered, and every seam closes.

Corollary 5.5 (elaboration soundness). For consistent  $e$  with  $\text{fee}(e) = k$ , the elaboration of Thm 3.7 inserts  $k$  coercions to yield  $e'$  with  $\vdash e' \checkmark$ ; by Thm 5.4,  $e'$  never gets convention-stuck. The  $k$  coercions are the disclosure normal form, and  $k$  is minimum.

Proof. Thm 3.7 produces coherent  $e'$ ; Thm 5.4 applies. ■

## 6.1 5.6 What Part II establishes; what is next

$\lambda_{\nabla}$  is sound: a program the type-checker accepts (consistent,  $\text{fee} = 0$ ) cannot, at run time, reach a convention mismatch — for any behavior of the undeclared (opaque) tools. The fee is the exact number of declarations needed to reach that guarantee. Convention coherence is thereby a genuine typing discipline, not a diagnostic: it prevents the failure rather than scoring its likelihood.

Part III admits feedback/recursion and internalizes the fee as a graded disclosure modality  $\square_r$  (grade = fee), whose two introduction forms are [CD] §7.0's revelation and contraction, connecting the grade to graded/quantitative type theory, and proves interface-evolution-as-

type-preservation ([CD] Conj 9.5). Part IV mechanizes Thm 5.4 (reusing [CD]’s Lean carriers and the latent-complexity measure  $\mu$  as the termination/measure scaffold) and extends Bulla into a prototype elaborator inserting disclosure-normal-form coercions on the MCP corpus.

Falsification. The soundness theorem fails if a coherent  $\lambda_{\nabla}$  term reaches `stuck_d` under some  $\varepsilon$  (refuting Thm 5.4) — which by Lemmas 5.1–5.2 would require  $R \not\subseteq \text{span } 0$  at  $\text{fee} = 0$  (a rank miscalculation) or a declared constraint violated at run time (a semantics error). Either is concrete and checkable, by exhibiting the term, the  $\varepsilon$ , and the reduction trace.

## 7 Part III. The Graded Disclosure Modality and Interface Evolution

Parts I–II treat coherence as binary:  $\text{fee} = 0$  or not. Part III internalizes the fee as a graded modality  $\square_r$ , with three payoffs. (1) Non-coherent terms become typeable, carrying their obligation count as a grade, so the type system measures incoherence rather than only rejecting it. (2) Recursion/feedback — excluded from the Part I core because it closes loops through application (breaking Prop 2.2) — is re-admitted, the grade carrying the feedback loop’s residual obstruction. (3) The deep result: an interface update that is chain-homotopy-trivial preserves the grade and transports the receipt, discharging the abstract invariance half of [CD]’s open Conjecture 9.5 in the  $\lambda_{\nabla}$  setting and yielding incremental recertification with a Lipschitz stability bound. The grade is a coeffect, but an unusual one: not a usage bound chosen by the programmer, but the witness rank forced by cohomology ([CD] Thm 5.1’s uniqueness). That is what makes  $\lambda_{\nabla}$  a graded type theory with a canonical grade.

## 8 6. The graded disclosure modality $\square_r$

### 8.1 6.1 The grade monoid

Grades live in the additive ordered monoid  $(\mathbb{N}, +, 0, \leq)$  of disclosure obligations. Additivity is forced by the program: the fee is additive over disjoint composition ([CD] Lemma 3.5) and sums along a tower ([HF] Paper I’s tower law), so composing two sub-systems adds their obligation counts.

### 8.2 6.2 The modality and its comonad structure

Definition 6.1.  $\square_r \tau$  is the type of a  $\tau$  carrying  $r$  unresolved disclosure obligations — a value of underlying type  $\tau$  whose convention coherence is obstructed by an  $r$ -dimensional witness class. Part I’s coherent fragment is the grade-0 fragment:  $\square_0 \tau \equiv \tau$ , and  $\Gamma \vdash e : \tau \checkmark$  (Part I) is exactly  $\Gamma \vdash e : \square_0 \tau$ .

Proposition 6.2 ( $\square$  is a graded comonad, with a debt-ordered subsumption). Over the additive monoid  $(\mathbb{N}, +, 0)$ , the family  $(\square_r)_{r \in \mathbb{N}}$  with counit  $\varepsilon : \square_0 \Rightarrow \text{Id}$  (extract a

coherent value), and  $\delta_{\{r,s\}} : \square_{\{r+s\}} \Rightarrow \square_r \circ \square_s$  (split obligations into an outer  $r$  and inner  $s$ )

is a graded comonad:  $\delta$  is coassociative and  $\varepsilon$  is a two-sided counit, the laws being exactly the associativity and unit laws of  $(\mathbb{N}, +, 0)$  (e.g.  $(\square_r \varepsilon) \circ \delta_{\{r,0\}} = \mathrm{id}_{\square_r}$  via  $r+0=r$ ). This is the additive instance of the coefficient-graded comonad of Gaboardi–Katsumata–Orchard [GKO] (cf. bounded linear logic’s  $!_r$ , [GSS]). Proof. The components are the canonical regroupings of obligation counts; each law reduces to the corresponding monoid identity in  $(\mathbb{N}, +, 0)$ . ■

Separately, the grade is an upper bound on obligations (a debt), so there is a sound subsumption coercion  $wk : \square_r \tau \rightarrow \square_s \tau$  for  $r \leq s$  — claiming more obligations than one has is always safe, costing only a redundant disclosure (A4b). Note the direction: this is the reverse of the usage-bound monotonicity of [GKO]/[GSS] (where  $\square_s \rightarrow \square_r$  for  $r \leq s$ ), because the  $\lambda_\nabla$  grade is a debt to be discharged, not a capability to be spent. We therefore keep  $wk$  as a subtyping coercion ( $\square_r <: \square_s, r \leq s$ ) rather than folding it into the comonad’s order structure; `reveal/contract` (§6.3), which lower the debt, are not part of the bare comonad since each requires a disclosure witness.

### 8.3 6.3 The two discharge operations (= [CD] §7.0)

A receipt is built from two structurally inequivalent operations — exactly [CD] §7.0’s revelation and contraction, now typed as the moves that drive the grade to 0:

- `reveal` (revelation, A4a):  $\mathrm{reveal} : \square_{\{r+1\}} \setminus, \tau \rightarrow \square_r \setminus, \tau$ , parameterized by an admissible disclosure of a class with nonzero cohomology — discharges one independent obstruction; grade  $-1$ .
- `contract` (contraction, A4b):  $\mathrm{contract} : \square_r \setminus, \tau \rightarrow \square_r \setminus, \tau$ , removes a redundant latent dimension (latent complexity  $\mu \rightarrow \mu-1$ ); grade unchanged.

These are not interchangeable: [CD] §7.0’s hybrid-semantics discovery is that revelation cannot witness A4b and contraction cannot witness A4a (verified there by `native_decide` counterexamples). A receipt is therefore a coKleisli morphism  $\square_r \tau \rightarrow \square_0 \tau$  — a sequence of exactly  $r$  reveals interleaved with any number of contracts, followed by  $\varepsilon$ . The cardinality  $r$  is forced (Thm 3.7); the `reveal/contract` mix is the free parameter ([CD] §7.0, §6.4).

Where `contract` bites. In the exact, single-dimension core of Parts I–II, `contract` is inert: DFD+CHP force the leaf-latent group  $C^\circ_{\text{seam}} = 0$  ([CD] Rmk 4.2.1), so every nontrivial disclosure has  $[\sigma] \neq 0$  and there are no coboundary classes to contract — receipts there are pure `reveal` sequences. `contract` becomes load-bearing only once feedback (§6.6) introduces leaf-latent state ( $C^\circ_{\text{seam}} \neq 0$ ), where a disclosure can re-expose inherited tool state (a coboundary) that A4b mandates leaves the grade unchanged. So the two moves are co-equal in general (the graded calculus with feedback), not in the recursion-free core.

## 8.4 6.4 Graded typing; the grade is the fee

The graded judgment  $\Gamma \vdash e : \square_r \tau$  assigns  $e$  the obligation count  $r$ . The principal (least) grade is the fee:

Proposition 6.3. The least  $r$  with  $\Gamma \vdash e : \square_r \tau$  derivable is  $\text{fee}(e)$ ; and  $r = 0$  iff  $e$  is coherent (Def 3.4). Proof. `reveal` lowers the grade by exactly one per discharged independent class and `fee` counts exactly those classes (Prop 3.5); `wk` can only raise it; so the least derivable grade is  $\text{fee}(e)$ , zero iff coherent. ■

## 8.5 6.5 Connection to graded / quantitative type theory

$\square_r$  is a graded comonadic coefficient (Petricek–Orchard–Mycroft [POM]; Brunel–Gabori–Mazza–Zdancewic [BGMZ]; Orchard–Liepelt–Eades [OLE]) in the additive ordered monoid of obstruction counts, and `reveal` is its graded-dereliction step (discharge one unit). Atkey’s Quantitative Type Theory [At] grades by a semiring of usages;  $\lambda_{\nabla}$  grades by obstruction count. The decisive difference — and the reason this is more than an instance — is that the  $\lambda_{\nabla}$  grade is canonical, not declared: in QTT/coefficient calculi the grade annotates how a programmer intends to use a variable; in  $\lambda_{\nabla}$  the grade  $\text{fee}(e)$  is computed from the composition and is the unique invariant respecting disclosure ([CD] Thm 5.1).  $\lambda_{\nabla}$  is thus a graded type theory whose grades are forced by a characterization theorem rather than chosen — a coefficient discovered in the program, not imposed on it.

## 8.6 6.6 Feedback and recursion, re-admitted

Part I’s recursion-free core made `reconcile` the sole loop-closer (Prop 2.2). The grade re-admits feedback. A feedback former `\mathsf{f}\{\text{loop}\}_d` feeds a value back through a self-map, closing a loop through application:

$$\frac{\Gamma \vdash e : \tau \rightarrow \tau \rightsquigarrow (O; R) \quad \ell := \text{fee of the closed loop on } d}{\Gamma \vdash \text{loop}_d(e) : \square_{\ell} \tau} \text{ (Loop)}$$

where  $\ell \in \{0, 1\}$  is the loop’s fee (0 if the self-map’s  $d$ -transport is declared and composes to the identity around the loop; 1 if opaque or unforced), generalizing to the fee of a mutual-feedback graph. So recursion is typeable, the grade tracking the feedback obstruction; Part I’s `reconcile` is the special case demanding  $\ell = 0$ . This restores Prop 2.2’s loop-closing role to the graded setting without sacrificing soundness (§6.7).

## 8.7 6.7 Graded soundness: the grade is the repair cost

The grade  $r = \text{fee}$  does not bound the number of seams that can fail — a single undeclared convention can make arbitrarily many reconciliations stuck at once (e.g.  $x_1 = \text{opaque}(x_0)$  reconciled with both  $x_0$  and a declared copy  $x_2 = x_0$  gives  $\text{fee} = 1$  yet two simultaneously-stuck seams). What  $r$  bounds is the repair cost: the number of disclosures that guarantee no seam

ever fails. One reveal discharges one independent obstruction class and thereby un-sticks every reconciliation that depended on it. This is the operational face of repair duality ([CD] Thm 6.2): obstructions, not failing seams, are what you pay to fix.

Theorem 6.4 (Graded Soundness — repair cost). Let  $\Gamma \vdash e : \square_r \tau$  be closed and consistent. Then:

1. (Progress)  $e$  reduces.
2. (Coherence floor) if  $r = 0$ ,  $e$  is never convention-stuck under any  $\varepsilon$  (this is Thm 5.4).
3. (Repair) the disclosure normal form (Thm 3.7) of  $r$  reveals drives  $e$  to a grade-0 term  $e'$  which, by (2), is never convention-stuck under any  $\varepsilon$ . Thus  $r$  is exactly the number of disclosures needed to guarantee no stuck seam, and a single reveal simultaneously certifies every reconciliation depending on the class it discharges.
4. (Tightness)  $r > 0$  iff some  $\varepsilon$  realizes a stuck reconciliation.

Proof. (1) is the progress argument of Thm 5.4. (2) is Thm 5.4. (3): by repair duality ([CD] Thm 6.2) and the disclosure normal form (Thm 3.7),  $r$  admissible reveals reach fee = 0 while preserving consistency (each pin is consistent with the forced frames), and Thm 5.4 then applies to  $e'$ ; one reveal discharges one class  $[\sigma_i] \in H^1$ , forcing every reconciliation whose constraint involves  $[\sigma_i]$  (Prop 3.5), of which there may be many. (4): an unforced class has a latent transport whose  $\varepsilon$ -instantiation can make a dependent reconciliation's holonomy nonzero (§2.5(b)); and  $r = 0$  is exactly (2). ■

So the grade is a cost, not a count:  $r$  disclosures suffice (and are necessary, by duality) to make a consistent term never-stuck, even when the number of seams those disclosures rescue is far larger. The computational QA (§5 and the Part III check) exhibits both faces: under-determined terms do get stuck (tightness, 9601/9601 in §5), and the number of stuck seams routinely exceeds  $r$  while  $r$  disclosures still repair all of them.

## 9 7. Interface evolution as graded type preservation (discharging the abstract half of [CD] Conjecture 9.5)

### 9.1 7.1 The problem

Tool manifests change: fields are added, conventions renamed, restrictions refined. Re-deriving coherence from scratch after every manifest update is the operational pain point ([CD] §9.5 calls this the most product-adjacent open problem). We want: an update that does not change the obstruction structure should preserve typing and transport the receipt — recertification for free — and a small update should change the obligation count by only a little.

## 9.2 7.2 Typed interface updates

Definition 7.1. A typed interface update  $u : \Gamma \Rightarrow \Gamma'$  is a morphism of convention sheaves — a natural transformation of the convention presheaves that respects the observable/latent split — inducing a chain map  $u^\bullet : C^\bullet(G_e) \rightarrow C^\bullet(G_{u \cdot e})$  on the seam cochain complexes ([CD] §3). The expanding updates — adding an observable field (extends 0) and renaming a convention (an isomorphism on a stalk) — induce forward chain maps. Not every update is forward. Convention narrowing (refining a restriction so a previously-accepted convention is rejected) has no forward map in the update category; following [CD]’s Phase-B analysis of Conj 9.5 it is modeled as a span  $\Gamma \leftarrow \Gamma'' \Rightarrow \Gamma'$ , and receipts do not transport across it — narrowing requires a fresh (partial) audit. This is the interface-update map [CD] Conj 9.5 left to be defined; the theorem below covers the forward, chain-homotopy-equivalent case, with narrowing carved out explicitly.

## 9.3 7.3 The preservation theorem

Theorem 7.2 (Graded Preservation under Interface Evolution; discharges the abstract invariance half of [CD] Conj 9.5 in  $\lambda_\nabla$ ). Let  $u : \Gamma \Rightarrow \Gamma'$  be a forward typed interface update (Def 7.1) whose induced chain map  $u^\bullet$  is a chain homotopy equivalence. Then for every well-typed  $e$ :

1.  $\text{fee}(u \cdot e) = \text{fee}(e)$  — the grade is preserved;
2. the obstruction class transports:  $u^* : H^1(G_e) \cong H^1(G_{u \cdot e})$  is an isomorphism;
3. any receipt (disclosure normal form, Thm 3.7) for  $e$  transports to a receipt for  $u \cdot e$ .

Hence  $\Gamma \vdash e : \square_r \tau$  implies  $\Gamma' \vdash u \cdot e : \square_r \tau$  with the receipt transported — recertification needs no re-derivation.

Proof. A chain homotopy equivalence induces isomorphisms on cohomology in every degree (standard homological algebra: if  $g^\bullet u^\bullet \simeq \text{id}$  and  $u^\bullet g^\bullet \simeq \text{id}$  via chain homotopies, then  $H^\bullet(u)$  is invertible). Hence  $\dim H^1(G_e) = \dim H^1(G_{u \cdot e})$  — (1) and (2). For (3) we argue at the level of cohomology classes, not pointwise contexts — a chain homotopy equivalence need not be a poset isomorphism and need not preserve the latent complexity  $\mu$ , so it cannot be assumed to carry an admissible disclosure to an admissible one verbatim. Instead: a receipt is a basis  $\{\sigma_i\}$  of  $H^1(G_e)$  with chosen admissible representatives (Thm 3.7); transport the classes by the isomorphism  $u^*$  to a basis  $\{u^*[\sigma_i]\}$  of  $H^1(G_{u \cdot e})$ ; then re-select an admissible representative of each  $u^*[\sigma_i]$  at a minimal latent context of  $G_{u \cdot e}$  (which exists by [CD] Lemma 5.3). The result is a valid receipt for  $u \cdot e$  of the same cardinality  $r$ , so the graded type  $\square_r \tau$  is preserved. (This is [CD]’s own Phase-B route for Conj 9.5: transport the class, re-realize the disclosure.) ■

[CD] §9.5 flagged this as “a routine specialization of standard cochain-homotopy invariance; the work is in defining the right notion of interface-update map.” Definition 7.1 supplies that map for the forward case; clause (3)’s class-level transport is the operational content turning invari-

ance into incremental recertification. We are deliberately modest about scope: this discharges the abstract invariance half of Conj 9.5 (the forward, chain-homotopy-equivalent updates); the narrowing half — where [CD]’s Phase B shows there is no forward map and receipts genuinely do not transport — is carved out in Def 7.1 and remains, with [CD], a non-transporting span requiring partial re-audit.

#### 9.4 7.4 Stability: bounded recertification under general updates

Most real updates are not chain-homotopy equivalences — they genuinely add or remove obstructions. For these we want a stability bound, so recertification is still local.

**Definition 7.3.** An update is  $k$ -local if it modifies the convention sheaf only on a sub-cover touching at most  $k$  convention generators (it adds or deletes at most  $k$  cochain dimensions).

**Theorem 7.4 (Stability / bounded recertification).** For a  $k$ -local update  $u$ ,

$$|fee(u \cdot e) - fee(e)| \leq k.$$

Consequently a coherent  $e$  remains coherent after at most  $k$  new disclosures, and the recertification cost is  $O(k \cdot |\text{update}|)$ , independent of  $|e|$ .

**Proof.**  $fee = \dim H^1$  is the rank of a quotient of cochain groups; adding or deleting  $k$  generators edits the coboundary by at most  $k$  rows / columns, and matrix rank is 1-Lipschitz under a rank-1 row / column edit, so  $\dim H^1$  changes by at most  $k$ . Equivalently, by excision / Mayer–Vietoris (axiom A6, [CD]),  $fee(u \cdot e) - fee(e)$  is determined by the change on the touched sub-cover and its overlap, bounded by its generator count  $k$ . ■

This is the persistence-stability analogue [CD] §9.5 points to (Cohen–Steiner–Edelsbrunner–Harer [CSEH]): the fee is 1-Lipschitz in the interface under generator-edit distance, just as a persistence barcode is stable in bottleneck distance under perturbation of the underlying data. It licenses a coherence-preserving versioning discipline: ship each manifest update with its  $k$ -local chain data; recertify by adjusting at most  $k$  disclosures, never re-auditing the whole composition. One caveat, inherited from §7.3: a convention-narrowing update (a tool tightening which conventions it accepts) is breaking — receipts do not transport across it — so the discipline carves narrowing out as the case that triggers a fresh, but still  $k$ -local (Thm 7.4), partial audit rather than a free transport. For the agent-systems setting this is the decisive practical property — tool manifests churn continuously, and Thm 7.4 bounds the audit cost of churn (forward or narrowing) by the size of the change, not the size of the system.

#### 9.5 7.5 What Part III establishes

The fee is realized as a graded comonadic modality  $\square_r$  (grade = witness rank), with revelation and contraction as its two discharge operations ([CD] §7.0) and a graded soundness theorem (Thm 6.4: the grade is the repair cost —  $r$  disclosures guarantee no stuck seam, and  $r = 0 \Rightarrow$

never stuck — not a bound on the number of seams that can fail, since one undeclared convention can break many seams at once and one reveal repairs them all). Recursion/feedback is readmitted, the grade carrying the loop obstruction. Interface evolution that is chain-homotopy-trivial preserves the grade and transports the receipt — discharging the abstract invariance half of [CD] Conjecture 9.5 in the  $\lambda_{\nabla}$  setting (with convention-narrowing handled as a non-transporting span, §7.3) — and general  $k$ -local updates change the grade by at most  $k$  (stability), giving incremental recertification bounded by update size. Part IV mechanizes Thms 5.4 and 7.2 (reusing [CD]’s Lean carriers, the  $\mu$ -induction, and the chain-homotopy machinery) and extends Bulla into a prototype elaborator that inserts disclosure-normal-form coercions and performs  $k$ -local recertification on the MCP corpus.

Falsification. Part III fails if a chain-homotopy-equivalent update changes the fee (refuting Thm 7.2 — which by cohomology-invariance would expose a definitional error in Def 7.1’s interface-update map), or if a  $k$ -local update changes the fee by more than  $k$  (refuting Thm 7.4). Both are concrete: exhibit the update, its chain data, and the before/after fees.

---

## 10 References

- [CD] J. Komkov, *A Witness Logic for Semantic Composition*, Res Agentica program, 2026.
- [LV] J. Komkov, *Local Validity Does Not Compose*, Res Agentica program, 2026.
- [PAR] J. Komkov, *Compositional Incoherence is a Parity*, Res Agentica program, 2026.
- [HF] J. Komkov, *The Coherence Fee and Hierarchical Decomposition of Coherence Fee*, Res Agentica program, 2026. (DFD+CHP sufficiency for the exact regime.)
- [SI] J. Komkov, *Signed-Incidence Structure in Compositional Verification*, Res Agentica program, 2026.
- [SCPI] J. Komkov, *SCPI: Predicate Invention Under Sheaf Constraints*, Res Agentica program, 2026.
- [1] C. A. R. Hoare, “An axiomatic basis for computer programming,” *CACM* 12(10) (1969), 576–580.
- [2] J. C. Reynolds, “Separation logic,” *LICS* 2002, 55–74.
- [3] R. Milner, “A theory of type polymorphism in programming,” *JCSS* 17(3) (1978), 348–375.
- [4] A. K. Wright and M. Felleisen, “A syntactic approach to type soundness,” *Information and Computation* 115(1) (1994), 38–94.
- [5] J. Hansen and R. Ghrist, “Toward a spectral theory of cellular sheaves,” *J. Appl. Comput. Topol.* 3(4) (2019), 315–358.

- [At] R. Atkey, “Syntax and semantics of quantitative type theory,” LICS 2018, 56–65.
- [GKO] M. Gaboardi, S. Katsumata, D. Orchard, F. Breuvar, T. Uustalu, “Combining effects and coeffects via grading,” ICFP 2016, 476–489.
- [POM] T. Petricek, D. Orchard, A. Mycroft, “Coeffects: a calculus of context-dependent computation,” ICFP 2014, 123–135.
- [BGMZ] A. Brunel, M. Gaboardi, D. Mazza, S. Zdancewic, “A core quantitative coeffect calculus,” ESOP 2014, LNCS 8410, 351–370.
- [OLE] D. Orchard, V.-B. Liepelt, H. Eades III, “Quantitative program reasoning with graded modal types,” Proc. ACM Program. Lang. 3 (ICFP) (2019), Article 110.
- [GSS] J.-Y. Girard, A. Scedrov, P. J. Scott, “Bounded linear logic: a modular approach to polynomial-time computability,” Theoretical Computer Science 97(1) (1992), 1–66.
- [CSEH] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, “Stability of persistence diagrams,” Discrete & Computational Geometry 37(1) (2007), 103–120.