

No Free Precedent

Typed Warrant Flow and the Yield of Institutional Judgment

John Komkov, Independent Researcher

July 2026

FRONTIER RECORD — ACCESSION GATE OPEN

Adjacent evidence, not an accession to the Technical Spine. Unmet accession gate: Foreign-authored applicability, independent adjudication, and non-captive compounding remain open.

LIFECYCLE technical-report | EVIDENCE formally-verified | RECORD /research/no-free-precedent
BUILT FROM research-status.yaml @ 2026-07-22

Abstract. Can one authorized judgment make future cases safely decidable? A label alone cannot. For any proper subset of labeled cases, two total decision functions can agree on every observed label and disagree on an unseen case. Reuse therefore needs a restriction: a reason predicate specifying which records the judgment governs. But a reason is not enough. Its adoption as precedent is an institutional act requiring authority, scope, prospectivity, conservativity, and preservation of prior refusals.

We implement this distinction in Bulla’s experimental Claim Flow profile. Four closed constructors convert evidence bundles into evidence claims, evidence claims into forum findings, final findings into reason-bearing precedent, and authorized premises into settlement claims. Every constructor carries a proposition-specific authority token and receipt. Case-only findings never generalize; persuasive findings never mutate semantic state; binding precedent applies only within its FRSL-1 reason, structured scope, authority regime, closure warrant, and semantic epoch. A fresh reason, wider scope, or new protected consequence is legislation.

We then freeze a 240-case, twelve-lineage planted-rule experiment with 48 untouched holdout cases. On this captive finite substrate, authorized precedent dominates exact-record preclusion on all twelve lineages, has median marginal yield nine, and reduces the median findings required for 80% eligible coverage by 75%, with zero observed unsafe or cross-boundary transfers. This is evidence that explicit reasons can amortize decisions when such reasons are available; it is not evidence that real institutions supply them, that FRSL-1 expresses them, or that the measured reuse has economic value. A 52-case laundering suite and precommitted derivation-budget receipts close the most immediate authority and resource attacks. Lean formalizes the core provenance and no-free-precedent results without **sorry**.

Internal research draft — Bulla Claim Flow v0.4

Evidence status: `INTERNAL_CAPTIVE`

Results are model-relative and not independently validated.

1. The proposition

The useful economic hope behind precedent is not that old outcomes resemble new outcomes. It is that an earlier authoritative act discharges part of the work needed to decide later cases. That hope is meaningful only if three quantities remain separate:

1. exact-case preclusion, where the same record is reused;
2. authorized reason-based precedent, where a stated predicate covers new records inside an adopted scope;
3. unauthorized generalization, where outcomes predict labels without an institutional rule.

Conflating these produces free precedent. A predictive model can find a stable pattern without possessing authority to make that pattern law. A forum can bind parties in one dispute without claiming a universal description of the world. A settlement can impose a consequence without proving its factual premises true. The protocol must therefore police conversions between claim kinds, not merely attach provenance to a final label.

Our question is deliberately narrower than a universal common-law compiler:

Given a finite record language, an explicit reason, and a valid precedential grant, how much future certified coverage can one adoption safely create without changing protected consequences or erasing refusals?

2. No free precedent

Let $\mathcal{C}$ be a finite or arbitrary set of cases, $\mathcal{S}$ a proper subset of cases with observed Boolean labels, and u an unseen case. Define two total decision functions. The first returns false everywhere. The second returns false on $\mathcal{S}$ and true outside $\mathcal{S}$. They agree on every observed label and disagree on u . Consequently, labels on a proper subset do not determine an unseen case without a hypothesis restriction.

The Lean theorem `no_free_precedent` constructs these extensions directly. It is intentionally elementary. Its role is constitutional, not mathematical spectacle: any engine that generalizes a case label must disclose the added restriction and the authority by which that restriction becomes operative.

An explicit reason predicate $r(c)$ supplies one such restriction. It still does not supply precedential authority. Claim Flow therefore treats reason extraction and precedent adoption as separate acts. The former can be computational; the latter is institutional.

3. Claim kinds and closed conversions

3.1 Claims

`WorldClaim` is a scope-relative assertion carrying an observation or domain warrant. `EntailmentClaim` is a conclusion inside a named model class. `EvidenceClaim` says that a record satisfies an authored evidence policy for a purpose. `InstitutionalFact` is a forum's finding for a bound case, purpose, scope, epoch, and appeal state. `PrecedentRule` is an explicit FRSL-1 reason with a declared effect. `SettlementClaim` is permission to impose a consequence.

These types do not form a ladder of truth. Entailment does not create an observation. A signature does not appraise evidence. A reserve does not make a premise true. A forum finding is not worldly truth. Settlement authority does not create precedential authority.

3.2 Authority

`ClaimFlowAuthority` holds four distinct grant collections: appraisal, adjudication, precedent, and settlement. Each `AuthorityToken` binds its exact permission, principal, authority-regime hash, scope hash, semantic epoch, and authorization-receipt hash. Constructors reject a token borrowed from a different permission, regime, scope, or epoch.

This is the semantic analogue of the confused-deputy problem: a component possessing one authority must not spend it as another. The mechanism builds on capability-security and proof-carrying-authorization ideas; the narrower Bulla contribution is proposition-specific authority for semantic and institutional claim conversion.

3.3 Constructors

The surface is closed:

```
APPRAISE(evidence bundle, appraisal token) -> EvidenceClaim
FORUM_FINDING(EvidenceClaim, forum token) -> InstitutionalFact
ADOPT_PRECEDENT(InstitutionalFact, reason, precedent token) -> PrecedentRule
SETTLE(EvidenceClaim | InstitutionalFact, settlement token) -> SettlementClaim
```

Every output embeds its authority-token and receipt hashes. Each edge is recorded in a one-action `ClaimFlowTrace`. Trace verification checks authority kind, receipt binding, epoch, topological order, and absence of output overwriting. Trace concatenation is valid only for the same action and epoch. This is not a global graph and supplies no global ordering of meaning.

4. Effects, refinement, and legislation

`CASE_ONLY` applies only to the source case or its canonical record equivalence. `PERSUASIVE` can be presented to an actor but cannot change the operative semantic surface. `BINDING_WITHIN_SCOPE` can enter precedent admission only after independent checks of:

- the final institutional-fact reference;
- the explicit canonical FRSL-1 reason and exact reason vocabulary;
- precedential authority;
- structured applicability scope;
- semantic epoch;
- finite conservativity;
- preservation of prior `RELY` and `REFUSE` cells.

The existing `build_precedent_admission` path now requires a successful `PrecedentAdoption`; a compiled J-tuple alone is insufficient. A same-epoch adoption may narrow a prior scope and retain a subset of protected consequences. Changing the reason, widening applicability, or adding a protected consequence emits `LEGISLATION_REQUIRED`. An epoch change is stale and requires supersession.

Historical decisions are append-only. A later review or remedy points back to the prior decision and records a new epoch; it cannot rewrite the prior status. This preserves prospectivity and accrued institutional facts without pretending that future governance is frozen.

5. Resource accountability and finality obstruction

Time limits are politically dangerous when treated as neutral. An interested party can choose a small budget to manufacture indeterminacy. Claim Flow makes an action-relevant resource exit valid only if a `DerivationBudgetPolicy` was authorized before execution and precommits logical limits for models, candidate atoms, formula nodes, branch nodes, and opposing pairs. It also pins permitted backend hashes and a challenge escalation.

`DerivationRunReceipt` records the problem and policy hashes, semantic epoch, backend and version hashes, logical counters, completed region hashes, search frontier, and output. Wall time and peak memory remain observational. A policy authorized at or after the run, an unpermitted backend, excess logical counter, or non-replaying frontier produces `UNAUTHORIZED_DERIVATION_BUDGET` or an invalid frontier.

The `explain-finality` command represents a finite finality policy as alternative sets of conditions. It enumerates missing conditions, removes set-dominated routes when enumeration completes, and emits a typed obstruction certificate. Incomplete enumeration yields a sufficient route with `minimality = UNRESOLVED`; it never claims impossibility. Multiple incomparable minimal routes yield `CHOICE_REQUIRED`. Burden coordinates remain separate, and reserve delta is `NOT_APPLICABLE` for reversible-only, human-review, and categorical-harm routes.

This mechanism uses the established idea of minimal unsatisfiable and correction subsets. Its distinct object is a typed blocker spanning semantic, grounding, authority, harm, epoch, recourse, and resource conditions.

6. Experiment

6.1 Corpus

The suite has twelve seam lineages—units, bounded time, interval boundaries, enums, null/absent, namespaces, integer rounding, delivery acceptance, evidence floors, revocation windows, authority scopes, and finite non-definability. Each lineage has twenty cases: sixteen design and four holdout. Every lineage binds a planted one-node FRSL-1 reason, authority regime, closure warrant, applicability scope, and epoch.

The sequences include base records, exact repeats, canonical equivalents, reason-preserving variations, irrelevant facts, fresh reasons, scope edges, authority changes, closure changes, conflict, resource exhaustion, and epoch supersession. Findings are labeled `CAPTIVE_MACHINE_DERIVED`. The corpus does not contain human adjudications.

6.2 Arms

Exact preclusion admits only known canonical records. Authorized precedent performs an explicit forum finding and binding adoption, then applies the reason only to eligible cases. The label-only learner predicts from outcomes but remains advisory and certifies no action. The declared-rule upper bound has the planted rule from the beginning. Both chronological and preregistered greedy design schedules are run.

The holdout is scored from the frozen design state without updating it. Thus a holdout result measures reuse created by design acts rather than a fresh finding on the held-out case.

6.3 Metrics and classification

The report retains `RELY`, `REFUSE`, and `ESCALATE` mass; findings; exact and reason reuse; typed exits; proof and formula size; logical counters; findings to 50%, 80%, and 90% coverage; and separate burden coordinates. Marginal yield is the number of additional future cases certified beyond the finding itself.

The preregistered `DEMONSTRATED_COMPOUNDING` threshold requires zero unsafe transfer, authorized dominance on at least nine holdout lineages, median yield at least two, and a median findings-to-80% reduction of at least 25%.

7. Results

The captive run cleared that threshold:

Measure	Result
Unsafe transfer	0
Cross-boundary reuse	0
Holdout lineages where authorized precedent beat exact preclusion	12/12
Median marginal precedent yield	9
Median reduction in findings to 80% eligible coverage	75%

This is the intended result of a planted-rule feasibility experiment. The authorized arm receives the explicit planted reason after the first eligible finding. The result therefore establishes that the constitutional gates do not destroy compounding when a compact valid reason exists. It does not measure the probability that a real opinion supplies such a reason or the cost of extracting it.

The label-only learner produced advisory predictions and no certified action. Authority, closure, conflict, fresh-reason, scope-widening, resource, and epoch cases remained typed exits rather than being counted as failures to generalize.

7.1 Denominator correction

The approved sprint plan described an existing 126-case partial regression cohort. The live frozen Golden v0.3 artifact contains 70. We retained those 70 as an instrumentation regression and recorded the mismatch as a falsified premise. We did not generate 56 cases and call them inherited evidence.

8. F11 laundering results

F11 contains thirteen archetypes in direct, nested, serialized-replay, and authority/epoch-mutation variants, for 52 cases. It attacks implicit world claims, missing appraisal, reserves treated as ground, forum-to-rule and adjudicator-to-legislator conversion, persuasive mutation, case-only and cross-scope reuse, stale precedent, categorical-harm pricing, settlement-to-truth, circular verification, and strategic resource starvation.

All cases receive their required typed rejection, legislation, staleness, or routing outcome. Six guard-removal mutant families are each distinguished by at least one F11 case. The standalone verifier recomputes case hashes, roots, holdout outcomes, classification metrics, regression bindings, self-hosting trace, and mutation witnesses under `python -I` without importing production Bulla. Clean-directory replay copies only the verifier and bound artifacts.

9. Formal results

The Lean bundle proves:

- authority inversion for every derived institutional claim;
- construction of disagreeing unseen-case extensions;
- non-generalization of case-only findings;
- reason-and-scope confinement of binding precedent;
- preservation of prior RELY and REFUSE under same-epoch refinement;
- preservation of historical decisions under supersession review;
- semantic-epoch invariance under larger derivation budgets;
- verification preservation under trace concatenation.

The bundle builds without `sorry`. The audit distinguishes these abstract theorems from the Python parser, receipt signatures, and serialization code, which are tested rather than proved in Lean.

10. Self-hosting

We replay the #171–#174 merge and freeze lineage as a local claim-flow corpus. Hosted checks supply an evidence substrate, maintainer merges are institutional acts inside the repository, and the freeze receipt and tag settle a versioned internal boundary. No `ADOPT_PRECEDENT` receipt was present, so the trace does not manufacture a reusable rule from CI success. It is classified `CAPTIVE_WITH_FOREIGN_SUBSTRATE`: GitHub ran foreign infrastructure but did not independently adjudicate semantic correctness.

11. Limits and next test

The crucial unknown has moved. Safe reuse is feasible when the reason is explicit, compact, authorized, and in-fragment. The next question is whether outsider-authored institutional materials naturally expose such reasons at a useful rate. That study must freeze the extraction method before collecting opinions, keep failed and ambiguous reason extractions visible, and distinguish the opinion’s stated ratio from a model’s post hoc rationalization.

Real institutions also introduce appeal, conflicting hierarchy, dicta, jurisdictional choice, strategic ambiguity, and contested record equivalence. The v0.4 calculus has typed places for these boundaries but no field evidence that its representation is accepted by participants. No production settlement, legal-validity, external-validation, or economic-value claim follows.

12. Conclusion

There is no free precedent. A past label does not decide a new case, and an adjudicator does not become a legislator by implication. Reuse begins only when an institution states or adopts a reason and authorizes its future effect. Bulla v0.4 makes that act explicit, scoped, conservative, prospectively receipted, and resource-accountable.

Inside a planted finite environment, those constraints still permit material compounding. That result is worth carrying forward precisely because its boundary is sharp: it demonstrates a mechanism, not a world.