

Supplementary Appendix B. Condensed staged run card for analyst–LLM use

B1. Standing rule and Context Package

Standing rules. Apply the L0–L5 layered conventions of Sections 2–3 throughout: do not treat a reported label, account, taxonomy element, disclosure line, amount, or glossary entry as an ontological primitive; do not attach L4 measurement to an unidentified LER or Institutional Act; do not let recognition, probability, valuation, or representation language constitute L1 existence; and preserve typed ambiguity where the source remains genuinely open.

Context Package. The Context Package is supplied to the LLM as a set of typed artifacts, not as free prose:

(i) OntoUML model — JSON serialization as the authoritative source of classes, stereotypes, generalizations, characterizations, and relation endpoints. The LLM reads class names and stereotype tags from this file rather than from natural-language summaries.

(ii) Reader-facing glossary — markdown giving for each class a short definition, layer assignment, “do not confuse with” notes, and execution role. The glossary explains terms but does not add ontological commitments absent from the model.

(iii) Executable kernel reference — the companion paper gives the L1 substrate definition: Performance LER, Holder Change LER, Price LER, Object LER, Holding LER, Reciprocity Relator, and the four-phase lifecycle. The LLM treats kernel definitions as given and does not redefine L1 terms.

(iv) Partition tables — Table 1 for reciprocity profiles and Table 2 for T0–T7 checkpoint questions and repair orientation.

(v) Bounded source artifact — the standards text, example, disclosure requirement, taxonomy element, reporting line, note, table, or other reporting surface being diagnosed, with explicit textual boundary.

(vi) Episode-specific prompt — the prompt selects the active subset of (i)–(v), states the immediate goal, and specifies the required output form.

Equivalent delivery through schema injection, retrieval-augmented prompting, or split context windows is compatible, provided each artifact is named and addressable in the prompt.

B2. Episode schema

For each episode, the prompt specifies the active Context Package, the immediate goal, ordered tasks, episode-specific constraints, and required output form.

B3. Episode 1 — Epistemic Scoper

Active Context: bounded source artifact, glossary support, procedure description, scoping rules.

Prompt: Given the bounded source artifact, identify the boundary, institutional frame, reporting frame, admitted evidence, and any additional evidence needed. Then perform an aggregation and boundary screen: identify possible netting, bundled LERs, omitted Other Parties, proxy classifications, reported residuals, internal splits, comparator-regime contamination, compressed phases, trigger shortcuts, temporal compression, service/holder-change collapse, and surface-only taxonomy distinctions. Do not let source headings or comparator regimes determine the reconstruction boundary. Do not perform formal OntoLER typing yet. Identify whether the bounded artifact is prescriptive or derived-reporting. Derived-reporting sections aggregate or summarize items whose L1 typing may be established elsewhere; do not diagnose them as defective merely because they do not reconstruct every underlying reciprocity.

B4. Episode 2 — Ontologist

Active Context: OntoLER classes and relations, glossary support, examples, partition tables, bounded source artifact.

Prompt: Using the bounded artifact, evidence, and prose draft, perform domain vocabulary control and reconstruct the L0 grounding basis and candidate L1 structure. Identify LERs, Institutional Acts, LER Change Classification, correlatives, reciprocity components, Triggering Situations, institutional effects, lifecycle phases, holder-change paths, service-manifestation paths, and relevant temporal structure such as due date, settlement horizon, or statutory term. Only after the L0–L1 structure has been reconstructed, attach L2–L5 overlays: business role, recognition assessment, measurement assessment, and representation scope. Where the source artifact is a disclosure, taxonomy element, note, table, line item, or other reporting surface, identify the representation anchors that preserve the lower-layer distinctions. Output a Typed Reconstruction.

Keep the LER partitions independent. Legal / Constructive concerns institutional basis. Validated / Hypothetical concerns whether a LER is sufficiently grounded as legal-economic reality. Executory / Conditional / Unconditional / Fulfilled concerns the lifecycle phase of an already reconstructed Performance LER. A legally instituted option or power may be Validated even though its exercise remains Conditional.

For each Reciprocity Relator, identify its components and profile. Do not infer reciprocity merely because correlatives exist. Correlativity pairs legal positions within a LER; reciprocity couples LERs, performances, compliance acts, allocations, objectified institutional positions, or authority-governed positions. Screen for bargaining, managerial/equity, and rationing reciprocity where relevant. If no L1 reciprocity is present, state whether the artifact is instead an object-lifecycle event, recognition act, measurement act, representation act, or other non-reciprocal institutional event.

If the clause concerns transfer, grant, retention, receipt, installation, delivery, or termination of an object-directed position, reconstruct the holder-change path explicitly at both levels: (a) name the persistent Object LER over the Underlying Object that survives the change; (b) name the prior Holding LER that terminates and the Holder that ceases to hold; (c) name the new Holding LER that is instantiated against the same Object LER and the Holder that begins to hold. Where the Object LER is itself an objectified Performance LER, such as an assigned receivable, transferable lease, bond, or share as institutional intangible, name the underlying Performance LER and state that the objectification has been instituted.

Specify Holder roles using the controlled vocabulary: Entity-as-Holder on the reporting side and Other-Party-as-Holder on the Other-Party side. Do not collapse a Holder Change into a generic “transfer” or “remeasurement” event.

B5. Episode 3 — Diagnostician

Active Context: Typed Reconstruction, checkpoint family T0–T7, relevant examples, selected formal constraints.

Prompt: Evaluate the Typed Reconstruction against T0–T7: grounding, LER integrity, reciprocity structure, lifecycle grounding, role-overlay attachment, recognition overlay, measurement overlay, and representation. For each checkpoint, state PASS, WARN, or FAIL; identify the failure signature; distinguish ontological priority from explanatory priority; and preserve typed ambiguity where the source remains genuinely open. T2 is conditional: absence of reciprocity is not a failure unless the bounded artifact presupposes reciprocity. Do not record T2 FAIL merely because a clause is methodological, presentational, measurement-oriented, or representation-oriented.

In addition to layer-specific failures, identify dependence-order inversions, where a reported L1 LER is constituted by L3 recognition, L4 measurement, or L5 representation inputs rather than by a constitutive L1 structure or triggering event. Before recording T6 FAIL, distinguish residual measurement from residual-LER compression. T6 fails only when monetary attachment precedes identification of the measured LER or Institutional Act, or its required specification / delimitation criterion. Before recording T7 FAIL, test whether the standard’s own required distinction cannot be reconstructed from the representation scope and its anchors; ordinary monadic line-item presentation is not itself a T7 defect.

If an apparent reciprocity component is actually an Object LER, reroute to the Ontologist as a typing problem unless the model explicitly treats it as an objectified Performance LER. The Object LER may be the object affected by holder change, but it is not silently assumed to be a reciprocity component.

Where a clause concerns a holder change, check that the two-level Object/Holding structure has been reconstructed: Object LER persistent, prior Holding LER terminated, and new Holding LER created. Conflation of the two levels — treating the Object LER itself as terminated and recreated — is a T1 typing failure, not a T3 trigger failure. Reroute to the Ontologist if the reconstruction has compressed the levels.

Where a clause concerns retained vs. transferred rights, such as leaseback, partial transfer, or sublease, the delimitation criterion required by T6 is the criterion that separates the retained Holding LER bracket from the transferred Holding LER bracket against the relevant Object LER structure. Monetary attachment without that criterion is T6 FAIL.

B6. Episode 4 — Repair Architect

Active Context: Diagnostic Profile, repair logic, relevant ontology elements, intended reader population.

Prompt: Using the Diagnostic Profile, propose the minimal useful repair and produce a Reconstruction Anchor for the intended reader population. State the relevant LER or Institutional Act, Other Party or reciprocity component where relevant, lifecycle status, triggering basis, holder-change or service-manifestation path where decisive, temporal qualifier if decisive, and any decisive overlay condition. Where the repair concerns representation, state the representation scope and the anchors needed to reconstruct the lower-layer distinction. State whether the repair concerns missing material at a layer, under-typed LER structure, missing specification / delimitation, measurement-before-delimitation, representation-anchor loss, or inversion of the dependence order.

B7. Re-entry rule

Return to Episode 1 for boundary, frame, evidence, or aggregation defects. Return to Episode 2 for LER, reciprocity, holder-change, service-manifestation, trigger, phase, temporal, overlay-attachment, or representation-scope defects. Return to Episode 3 when a repair is inadequate and must be retested before retyping.

B8. Human / LLM division of labor

The LLM supports lexical screening, candidate boundary suggestions, provisional prose drafts, draft L0–L5 reconstructions, draft representation-anchor checks, and draft checkpoint scorecards. The analyst retains responsibility for source selection, scoping correction, acceptance or revision of typing, adjudication of ambiguity, and final diagnosis and repair. The LLM treats the On-toUML model and glossary as the controlled vocabulary: it does not introduce class names absent from the JSON, does not redefine terms whose definition the glossary supplies, and flags terminological gaps as Typed Ambiguity rather than filling them silently.