LC

Finance Agent Integrity Pack

A practical review pack for AI agents that touch financial models.
Plausible is not enough. The work has to be inspectable.

10-Point Integrity Checklist

1. The original instruction is recorded.
2. Every changed assumption has a name, unit, period, value, and source type.
3. Important unchanged assumptions are explicitly listed.
4. Changed formulas or dependencies are named.
5. Base case and scenario branch are separated.
6. Recalculation evidence exists before output readback.
7. Validation checks show expected result, actual result, and pass/fail.
8. Failed, skipped, or unavailable checks are disclosed.
9. The durable model state or artifact path can be opened by a reviewer.
10. The final answer points to evidence, not only prose.

Failure Modes Covered

Failure mode	Control question
Hidden Assumption	Can every material input be found in the assumption register?
Base Case Overwrite	Can base and scenario outputs be reconstructed side by side?
Scenario Contamination	Did a scenario assumption leak into the base case?
Formula Drift	Did an input-only edit leave calculation logic unchanged?
Timeline Misalignment	Were annual, monthly, and partial-period assumptions converted correctly?
Cash Blindness	Did the agent show the cash consequence of changed payment timing?
Phantom Precision	Does reported precision match the quality of the underlying assumptions?
Validation Theatre	Are validation checks exposed as inspectable rows, not only claimed?
Export Amnesia	Can the spreadsheet export be tied back to durable model state, source revision, hashes, and validation?
Source Confusion	Are public, synthetic, redacted, and private inputs labeled?
Stale Formula Readback	Is there proof that dependent outputs recalculated after the edit?

Failure mode	Control question
Workbook Round-Trip Damage	Did formulas, validations, named ranges, and hidden audit sheets survive export/import?
Missing Recalculation Proof	Does the packet expose a recalc engine, timestamp, dependency output, and pass/fail certificate?
Skill Packaging Drift	Do manifest, registry, examples, changelog, advertised controls, and smoke-test evidence match?
Unsafe Permission Carryover	Do active permissions and inputs stay inside the declared source and permission boundary?
Stale Scenario Branch	Does the scenario reflect the current parent model version?
Proofless Tool Success	Can the downstream artifact be opened after a successful tool call?

Minimum Reviewer Packet

Instruction log, assumption register, changed calculation list, scenario diff, validation rows, unresolved limitations, durable model or artifact link, and export only as a downstream artifact.

Public Proof

Failure modes: https://lechiffre.cc/finance-agent-failure-modes/
Readback proof: https://lechiffre.cc/finance-model-readback-proof/
Public synthetic Layerz model: https://layerz.cc/models/99f43c06-d332-4bd7-ac66-b531eeb351ff

LeChiffre AI. Built with Layerz where noted. Educational modeling and decision-support material, not financial, accounting, tax, legal, or investment advice.