METHODOLOGY

Kodak

Retrodiction methodology

Paul Campillo

Version 1.0 · March 2026 · Complete

What a retrodiction is

A retrodiction applies a diagnostic framework to a historical case where the outcome is already known—but runs the analysis as if standing at a past date, using only evidence available at that date. It then scores the diagnosis against what actually happened.

The purpose isn't to explain the past (any narrative can do that after the fact). The purpose is to test whether the framework could have **predicted** the outcome from the evidence available **before** the outcome was known. If it can, the framework captures real structure. If it can't, the framework is just storytelling.

Why Kodak

Kodak is the ideal retrodiction subject for three reasons:

The outcome is unambiguous. Chapter 11 bankruptcy, January 19, 2012. Revenue declined from $16B to $6.1B. Market cap from $31B to near-zero. Workforce from 145,000 to fewer than 8,000. No ambiguity about what happened.

The conventional narrative is well-established. "Kodak failed to adapt to digital photography." This gives us a benchmark: Genesis must produce something the conventional narrative cannot—more specific, more mechanistic, more predictive.

The data is rich and public. SEC filings, market data, industry reports, analyst coverage, academic research, and management statements span the full trajectory. No evidence gaps.

The protocol sequence

Genesis ran four protocols in sequence, each consuming the prior one's output:

STEP	PROTOCOL	WHAT IT DOES	OUTPUT
1	Research	Collects evidence at each checkpoint, admitting only information available at that date	Evidence packages: T0 (104 claims), T1 (24 claims), T2 (21 claims)
2	Diagnostic	Scores the system's capacity, market access, margin, navigation, and trajectory at each checkpoint	SIRF scores, margin, posture, boundary inventory, bet register at each cutoff
3	Prediction	Generates counterfactual forecasts—"if standing at this date, what would Genesis predict?"	Scenario bundles, trajectory projections, 19 distinctive predictions
4	Retrodiction	Scores all diagnoses and predictions against the known outcome	Brier scores, directional accuracy, error analysis, pattern library entry

The complete scored record is hash-locked in the Seer forecast registry as SEER-2026-0002.

The three checkpoints

CHECKPOINT	DATE	WHY THIS DATE	EVIDENCE CLAIMS
T0	December 1996	Peak—highest revenue, stock, brand value. Maximum distance from outcome. Hardest test.	104
T1	December 2003	Midstream—film declining, digital exploding, digital pivot announced. Tests whether Genesis tracks the transition correctly.	24 (delta scan)
T2	December 2007	Late—crisis forming but not yet obvious to all observers. Last checkpoint before outcome becomes visible.	21 (delta scan)

Admissibility discipline

The cardinal rule of retrodiction: **no future information contaminates past analysis.**

At each checkpoint, the diagnostic used only evidence that was publicly available by that date. No analyst could have known in 1996 that Kodak would file bankruptcy in 2012. The Genesis diagnostic didn't know either—it diagnosed from the evidence at hand and let the framework produce the forecast.

Specific safeguards:

- Every evidence claim carries an `observed_date` tag and an admissibility tier (T1 = available at checkpoint; T2 = available by next checkpoint; T3 = inference from available evidence)

- No claim references events after the checkpoint date
- Diagnostic scores reference claim IDs—no invented evidence
- Predictions are timeboxed to horizons forward from the checkpoint, not backward from the known outcome
- All artifacts frozen before scoring begins—no post-hoc adjustment

Scoring methodology

Two levels of scoring, measuring different things:

Bet-level Brier score: Scores each individual prediction against its binary outcome. For each prediction with assigned probability p and outcome o (1 = correct, 0 = incorrect):

$$\text{BS} = (p - o)^2$$

Aggregate bet-level Brier = mean across all 17 predictions. Result: **0.050**. Perfect calibration = 0.0. Random guessing at 50% = 0.25.

Scenario-level Brier score: Scores the primary scenario probability against the aggregate forecast accuracy. The base scenario was assigned 55% probability; the aggregate forecast score (weighted component accuracy) was 0.978. Result: **0.183**.

The gap between these scores is informative. A bet-level Brier of 0.050 with a scenario-level Brier of 0.183 means the individual structural predictions were highly accurate but the overall scenario confidence (55%) was conservative relative to how right the diagnosis turned out to be. For a first retrodiction with unvalidated methodology, that conservatism was appropriate.

The Seer registry (SEER-2026-0002) stores the scenario-level score (0.183) alongside the forecast score (0.978) and differential accuracy (+0.848 vs. baseline frameworks).

Directional accuracy: Did the prediction correctly identify the direction of change? Binary—correct or incorrect. Partial credit (0.7–0.8) awarded when direction was correct but threshold was slightly missed. Result: **17 of 17 (100%)**.

Benchmark comparison: Each Genesis prediction compared against what Wall Street consensus, McKinsey/Porter competitive analysis, and Christensen disruption theory would have predicted at the same checkpoint with the same evidence. Differential accuracy: +0.848—Genesis added substantial predictive value beyond all three baselines.

What the retrodiction does not claim

This retrodiction demonstrates that Genesis captures structural patterns that conventional analysis misses. It does **not** claim:

- **Certainty.** The T0 predictions carried 55–80% confidence—appropriately uncertain at peak. Genesis doesn't claim to have known Kodak would fail. It claims to have identified **why** it was structurally vulnerable, and assigned probabilities that proved well-calibrated.
- **Uniqueness.** Other frameworks (particularly Christensen's disruption theory) identified Kodak as at-risk from digital photography. Genesis's advantage isn't the risk identification—it's the **mechanism specificity**: which internal channel was broken, why each strategic pivot would fail for the same reason, and the timeline of margin consumption.
- **Hindsight elimination.** Despite rigorous admissibility discipline, the analyst (running the retrodiction in 2026) knew the outcome. Unconscious bias toward the known result is possible. The safeguards (evidence dating, claim ID tracing, frozen artifacts) mitigate but cannot fully eliminate this risk.

The honest claim: Genesis identified structural vulnerability, specific failure mechanisms, and approximate timelines from peak-era evidence, with a bet-level Brier of 0.050 (five times better than target) and 100% directional accuracy. The framework appears to capture real structure in organizational transformation dynamics.

The live test is the Tesla prediction—where hindsight contamination is impossible because the outcome hasn't happened yet.

Protocol: Retrodiction Protocol v2.3 | Diagnostic Protocol v3.3 | Prediction Protocol v2.3 | Research Protocol v2.7 | **Evidence:** 149 evidence claims across three cutoffs. Public filings, market data, academic sources, industry reports.

Forecast registered: SEER-2026-0002 · github.com/wakeupsmiling/seer · Bet-level Brier: 0.050. Scenario-level Brier: 0.183. Full analysis: genesistheory.org/kodak