

OPERATING CASE STUDY / CUSTOMER EXPERIENCE & LOGISTICS

Late-Delivery Recovery & Action Routing

Turn a disputed delivery label into a verifiable operating decision

An anonymized operating case on diagnosing major late-delivery categories, reconciling competing data views and assigning accountable follow-up. It separates a recorded location from current custody, cause and preventable customer harm.

PROJECT CONTEXT / DELIVERY DIAGNOSIS & ACCOUNTABLE EXECUTION

BLUF / BOTTOM LINE UP FRONT

Route investigation through verified evidence.
Keep labels separate from root cause.

ANALYTICAL FINDING

Matching dashboard totals do not show which unclassified cases another field explains. A linked cross-tab and aligned event clocks are needed before converting a label into an owner or recovery claim.

CONTRIBUTION IN THE SOURCE RECORD

Commissioned and adjudicated the analysis, reconciled versions and preserved retractions. Provisional registry hoop support.

EVIDENCE STATUS

Originals unrecovered. Operating fixes, mature causal outcomes and realized financial benefit remain unestablished.

BLUF: analytical inference. Source contribution: historical record. Client and internal identities anonymized.

EXECUTIVE ABSTRACT / DECISION AND OWNERSHIP

Commission the linked diagnosis before transferring responsibility

[AI/PD] The proposed Attribution-to-Action Contract (AAC) makes each handoff explicit: label, observed event, current custody, actionable cause, accepted owner and mature outcome. A persuasive category name supplies none of the later evidence by itself.

Abstract element	Decision-relevant synthesis
Context [HR]	A late-delivery action plan challenged legacy category names, incomplete reason codes and conflicting data views. The author commissioned analytical runs, reconciled their disagreements and assembled executive and operating artifacts.
Objective [PD]	Direct scarce operational attention to independently verified causes while preserving customer promises. Distinguish diagnostic routing from a formal ownership decision, an implemented remedy and a realized benefit.
Approach [RC/AI]	Reconcile matched coverage, category families, shipment exposure, calendar definitions and competing source versions. Test the inference from field marginals to missing reasons and from last recorded stage to current custody.
Findings [AI]	Matching totals do not identify the missing-reason subset. Recorded stage supports a location-based investigation, not complete custody or control. A decline inside one outcome family can coexist with unchanged total affected units.
Decision value [PD]	Release a bounded linked-record review and an owner acceptance protocol. Correct affected artifacts; preserve unresolved categories and competing versions. Evaluate a defined response package at fixed capacity on original eligible episodes.
Implications [PD]	Data certifies the population and event bridge. Operations verifies current state and accepts actions. Quality adjudicates outcomes; Finance prices incremental work and mature customer effects. Taxonomy changes receive no automatic savings credit.

EXECUTIVE DECISION

[PD] Approve the evidence and ownership gate. Require linked entities, valid event chronology, coverage accounting and a named action owner before recommending recovery. Keep existing customer-entitlement rules in force throughout the proposed evaluation.

Basis: selected project register S1. [HR] historical record; [RC] recalculated; [AI] analytical inference; [PD] proposed design. This is an analytical and action-planning case, not evidence of implemented operating benefit.

EXHIBIT 1 / ATTRIBUTION-TO-ACTION CONTRACT

A last recorded stage narrows the search without proving the cause

[HR] The source obtained a current-period stage drill for two disputed delivery categories and challenged the implication of their names. [AI] That is useful diagnostic evidence. It does not establish where every parcel was at a later decision, which process caused the failure or which team could prevent it.

OBSERVED LOCATION IS A STARTING POINT FOR INVESTIGATION

Proposed Attribution-to-Action Contract / each transition requires new evidence



A recorded stage is neither continuous custody nor a preventable cause.
A changed label is neither an implemented remedy nor a changed customer outcome.

Basis: original AAC synthesis of the source diagnostic and unresolved ownership handoffs.

Evidence state	Decision it can support
Category and event [HR/AI]	Describe the recorded outcome and the last captured stage within a declared population. A legacy label can be misleading; an observed scan can be stale, incomplete or followed by an unrecorded transition.
Current custody and cause [PD]	Reconcile event, posting, promise and decision clocks. Verify the physical state and the actionable mechanism with independent evidence; distinguish upstream causes from the location where consequences surface.
Owner and remedy [PD]	Assign an interim investigation lead, then obtain explicit acceptance of a defined corrective action. State dependencies, authority, completion evidence and a reopening trigger.
Mature outcome [PD]	Measure receipt, lateness, nonreceipt and customer burden on original eligible episodes. A renamed category or reassigned ticket does not establish a changed physical outcome.

THE RIGHT USE OF LOCATION EVIDENCE

[AI/PD] Start investigation at the last supported location. Escalate on a verified cause and an accepted action contract. Do not turn the location of a consequence into an automatic assignment of blame or financial accountability.

Basis: S1 tier-stage analysis, ownership decisions and unresolved follow-up. AAC is a proposed framework; source terminology is anonymized.


EXHIBIT 2 / COVERAGE AND MISSINGNESS

Preserve the unmatched population when interpreting a dominant stage

[HR] The stage drill covers different proportions of the two source categories. One has a substantial unmatched portion. [AI] A dominant stage among matched units cannot be applied directly to the entire category; missingness may reflect load lag, linkage failure or a different operating path.

MISSING UNITS TURN A POINT ESTIMATE INTO A RANGE

Hypothetical: 75% matched coverage; 96% in-station among matched; unmatched state unknown



Whole-population share = coverage × matched share + (1 – coverage) × unmatched share.

Basis: algebraic missing-data bound. Wholly hypothetical; no source coverage or stage share disclosed.

[RC/AI] Let c be the matched fraction, p the in-station share among matched units and q the unobserved in-station share among unmatched units. The whole-population share is $\theta=cp+(1-c)q$, bounded by cp and $cp+1-c$ when q lies between zero and one. The illustration changes q while holding the observed matched portion fixed.

Question	Required evidence [PD]
Coverage mechanism	Account for every unmatched identifier and distinguish late loading from true non-attribution. Test differences by node type, process, calendar and outcome maturity.
Time of observation	Determine whether the stage describes last scan, current custody or a later inferred state. Require evidence at the action decision rather than treating the last event as continuously true.
Operating responsibility	Trace an actionable mechanism and obtain owner acceptance. Even a robust majority location finding does not establish preventability or exclude upstream contributions.

A BOUND CAN SUPPORT A SEARCH PRIORITY

[AI] If the lower bound still indicates a majority, that justifies prioritizing the location in a diagnostic plan. It does not license a whole-category causal estimate, a recovery commitment or a change to customer-entitlement timing.

Basis: S1 matched-versus-category coverage. Diagram is wholly hypothetical; confidential matched shares, coverage anchors and counts are retained in the private annex.

EXHIBIT 3 / MARGINALS AND LINKED RECORDS

Matching totals do not reveal which cases explain the missing-reason group

[HR] Two dashboard views reported the same category total while using different fields: a reason code and a last status. The source used the second view to interpret a large unclassified reason segment. [AI] Equal marginal totals do not establish identical entities, and identical entities alone do not identify their joint classification.

IDENTICAL MARGINALS CAN SUPPORT OPPOSITE SUBSET STORIES

Hypothetical aligned population: 100 entities; unknown reason U = 40; selected status S = 50

CROSS-TAB A / NO UNKNOWN CASE HAS S			CROSS-TAB B / ALL UNKNOWN CASES HAVE S		
Reason group	S	Not S	Reason group	S	Not S
Unknown U	0	40	Unknown U	40	0
Known reason	50	10	Known reason	10	50

Totals alone cannot choose between these tables. Obtain the entity-level join.

Basis: exact synthetic cross-tab construction. Same entities is an assumption, not a consequence of tied totals.

[RC/AI] Conditional on one aligned population, if U denotes an unknown reason and S a selected status, their overlap satisfies $\max(0, P(U)+P(S)-1) \leq P(U \cap S) \leq \min(P(U), P(S))$. Many cross-tabs can have the same marginals. The synthetic examples show opposite explanations of U with unchanged totals.

Release gate	What it resolves [PD]
Entity and scope join	Confirm the same identifiers, eligible weights, country/node scope, date basis and versions. A coincident total is an accounting check, not proof of entity equality.
Reason × status cross-tab	Join the fields on those entities and preserve null, unknown and conflicting states. Inspect within-group composition; do not apply a whole-category status share to the unknown-reason subset.
Independent causal sample	Review the event history and physical evidence for a stratified sample. A linked status explains an administrative pattern only; an actionable cause and responsible owner still need verification.

KEEP THE UNRESOLVED SEGMENT OPEN

[AI/PD] Use the alternate field to formulate a testable explanation. Assign an interim diagnostic owner and a dated evidence request. Do not publish a settled root cause from two separate marginal distributions.

Basis: S1 reason/status views and contested category ownership. Hypothetical table illustrates partial identification; source overlap bounds are conditional and private.


EXHIBIT 4 / FAMILY RECONCILIATION AND EXPOSURE

A smaller subcategory can conceal unchanged total affected volume

[HR] The author adopted a rule to sum every sibling outcome family before reporting a subcategory change. The register also preserves conflicting totals, rates and extracts. [AI] The rule is valuable only when the families are exhaustive, nonoverlapping and evaluated on the same eligible, mature cohort.

A FAMILY DECLINE CAN BE OFFSET WITHOUT ANY REDUCTION IN TOTAL

Wholly hypothetical fixed affected population / four exhaustive, disjoint outcome families



[RC/AI] In the hypothetical illustration, a visible family falls from 60 to 45 while the other families absorb the difference; total affected units remain 100. If shipped exposure independently rises from 100 to 110, the pooled rate index falls to 90.91 despite no reduction in affected units. Neither movement identifies treatment benefit.

Reconciliation	Decision consequence
Conservation [PD]	Sum unit-grain outcomes across the full family set. Test duplicates, exclusions and transitions. An offset in aggregate categories does not prove that the same parcels migrated.
Rate arithmetic [PD]	Recompute each rate from its declared numerator and exposure. Keep count, weighted units and rate separate. Do not average incompatible dashboard and family-sum values.
Vintage and maturity [PD]	Align calendar boundaries and freeze a mature outcome window. Compare complete source versions separately; an under-matured current period cannot support a banked improvement.

MEASUREMENT RULE BEFORE PERFORMANCE CLAIM

[AI] The source supports a correction to reporting discipline. It does not support selecting one of its inconsistent snapshots as a realized improvement. Preserve the competing reads until an event-level bridge reconciles them.

Basis: S1 net-across-families rule, shipment exposure and source-version conflicts. Every number in the exhibit is hypothetical; source arithmetic is documented privately.


EXHIBIT 5 / QUERY AND OBSERVATION VALIDITY

A missing row in a returned view is not proof that the data cannot exist

[HR] An earlier source artifact claimed two categories were absent from a view by construction. A later browser export showed them. The correction was recorded, but repair of every shipped artifact is not established. [AI] The failure was a jump from incomplete observation to a claim about the system’s structure.

RETRIEVAL FAILURE AND STRUCTURAL ABSENCE ARE DIFFERENT CLAIMS

Source analytical correction / every interface needs an explicit population and completeness check



Basis: S1 incomplete returned view, later browser evidence and unresolved repair of shipped materials.

Source trap	Validation before interpretation [PD]
Silent filter or cap	Verify applied filters, returned row limits and population totals against an independent path. Identical results after a filter change can indicate an ignored filter, not stable performance.
Different field meanings	Reason, last status, process stage and outcome family answer different questions. Preserve field names, derivation logic and taxonomy versions privately; do not equate matching words across fields.
Different calendar grains	Align event time, posting time and week boundaries. A join that matches a week label can still pair different days, vintages or eligible node populations.
Incomplete or stale export	State retrieval status and source authority. Re-run decisive queries after access returns. More expansive coverage from a weaker verification path does not supersede a narrower reproducible result.

RETRACTION MUST REACH THE ARTIFACT

[PD] Record the corrected claim, affected artifacts, accountable editor and replacement version. Confirm that the executive document and call materials were replaced. Logging a correction does not establish that decision makers stopped using the old statement.

Basis: S1 reversal of the structural absence claim, calendar/scope pitfalls and incomplete retrieval. The portfolio case retains the correction and does not repeat the disproven claim.


EXHIBIT 6 / ANALYTICAL CONTRIBUTION AND CHALLENGE

Credit the author for adjudication without converting review volume into proof

[HR] The author commissioned agent-assisted analyses, steered the questions, independently checked a headline drill, reconciled competing runs and retained explicit retractions. [AI] The source also seeded a preferred interpretation. Challenge review must permit it to fail; agreement among agents cannot replace original evidence or a valid inference.

ANALYTICAL LEADERSHIP INCLUDES CORRECTING THE PREFERRED STORY

Historical contribution is separated from commissioned analytical work and operating implementation



Contribution	Process-tracing grade and boundary
Framing and integration [HR/AI]	Provisional registry hoop support: set the output contract, steered the analysis, selected evidence and integrated operating artifacts. Analytical labor included commissioned agents; no direct-report team contribution is invented.
Challenge and restraint [HR/AI]	Provisional hoop support for re-derivation, conflicting-version reconciliation and withdrawal of unsupported claims. A recorded correction is not a quantified historical avoided loss.
Current source replication [RC]	This case checks the supplied register and reconstructible arithmetic. Original-source recovery determines whether stronger provenance can be claimed; prior assertions of verification are not inherited as fresh checks.
Operating implementation [HR/AI]	Unestablished. Data enablement, formal taxonomy changes, local remedies and customer-policy changes depended on other owners. Producing an action plan is not evidence that those fixes shipped.

THE DEFENSIBLE PORTFOLIO CONTRIBUTION

[AI] The case demonstrates analytical leadership through a disputed operational decision: preserving useful findings, retracting unsound inferences and defining what evidence the next owner needs. It claims no independent external review or realized savings.


Basis: S1 personal contribution, commissioned analytical work, retractions and implementation boundaries. Stronger process-tracing grades require recovered, discriminating evidence.

EXHIBIT 7 / ACTION OWNERSHIP AND CUSTOMER OUTCOMES

Close the ownership handoff with a verified remedy and a later outcome

[HR] The source named fixes and owners, but durable changes remained unresolved in the supplied record. [PD] Convert each diagnostic item into an accepted action contract with a specified population, evidence requirement, authority, due date and outcome readback.

EVERY ACTION NEEDS ACCEPTANCE, EVIDENCE AND AN OUTCOME READBACK
Proposed action contract / scope, authority, dependency, due date, evidence and reopening trigger



Accountable function [PD]	Acceptance evidence and boundary [PD]
Data and analytics	Certify linked entities, family conservation, coverage, event clocks and versioned definitions. An access request or query draft does not establish data delivery.
Operations and Quality	Verify the present exception, identify an actionable cause and authorize the response through the established operating channel. Retain unresolved cases and independent outcome adjudication.
Policy and program owners	Approve taxonomy or workflow changes separately. Preserve customer promises and existing entitlement timing during the proposed test; the historical refund-hold thesis belongs to a separate unresolved project.
Finance and evaluator	Compare response policies at equal capacity on original eligible episodes. Value mature receipt states and incremental effort. A label-only transfer earns no physical-benefit credit.

RELEASE BY EVIDENCE, ACCEPTANCE AND OUTCOME

[PD] Proceed from a bounded linked-record diagnosis to a defined response pilot only when ownership and measurement are ready. Scale requires useful net value and customer safeguards. A closed ticket or cleaner taxonomy is an intermediate result.

Basis: S1 unresolved owner handoffs. Proposed AAC release gates and later randomized response model; no historical deployment or customer-benefit claim.


TECHNICAL ANALYSIS

Connect an attributed record to an authorized, testable response

[AI/PD] The proposed Attribution-to-Action Contract (AAC) separates recorded location, current custody, actionable cause, accepted responsibility and verified response. A convincing location analysis does not identify every later transition.

ATTRIBUTION IS AN INPUT TO ACTION, NOT PROOF OF CONTROLLABILITY

Proposed AAC causal structure / physical outcomes and administrative labels remain separate



Primary ITT: physical burden from Y plus W; compare its reduction with K. B is secondary.

All original eligible episodes remain. P is information used by E, not a validated cause of Y.

[PD] Z assigns the evidence-and-routing offer; E is usable evidence, R the accepted route and T the delivered response. Y is physical disposition. R and T incur variable work W; Z commits fixed coordination cost K. The proposed primary contrast includes physical burden and W, with K charged separately.

[AI/PD] Physical history H and recording quality V determine last-recorded stage P. P informs E; it does not itself cause or certify recoverability. Y and V affect administrative family B. Independent reference quality J and Y determine observed disposition O. C/U includes severity, actionable cause, capacity, business line and concurrent work.

[AI/PD] The historical register concerns evidence and analytical decisions, not randomized Z. Estimate a total offer effect across all original eligible episodes. Conditioning on accepted routes, completed responses or later concession labels would select mediators or outcomes. Shared recovery capacity must remain within the assignment domain.

Basis: registered analytical paths [HR]; proposed AAC causal structure [AI/PD]. Arrows are hypotheses; proxy location is not an instrument or proof of causal ownership.

TECHNICAL ANALYSIS

Preserve the net family count without discarding outcome severity

[HR/AI] The source’s instruction to reconcile sibling late-delivery families protects against a narrow reporting win. Its numerical examples still contain population, count, weighting and version forks that require separate reproduction.

Quantity	What the comparison can establish
Net administrative families [AI]	A sum is meaningful only for a mutually exclusive, exhaustive partition of the same unique units, period, exposure and weighting rule. An unchanged sum cannot prove the same entities migrated.
Physical disposition [PD]	On-time receipt, late receipt and nonreceipt by a frozen horizon are distinct outcomes. A recovery from nonreceipt to late receipt can reduce severity even if the combined late/nonreceipt share is unchanged.
Resource burden [PD]	Apply explicit weights to physical outcomes and separately count additional routing/recovery work. A classification correction alone has no physical outcome credit.

[AI/RC] For a fixed unit universe, pure reassignment among family labels conserves the union count. With common exposure E, the pooled rate is the union count divided by E. Exposure growth can lower that rate without fewer affected units. Separately normalized family indices cannot be added as though they were counts.

[HR/AI] An overall total equals the sum of two country subtotals exactly. That suggests a scope explanation; matching entities, periods and weights must be verified before the apparent discrepancy is closed. Other quoted count, rate and net-change branches remain unreconciled; no preferred branch earns priority because it tells a cleaner story.

[PD] Positive benefit is mean burden under usual response minus mean burden under the offer. Report each physical-outcome contrast and its uncertainty beside the composite. Freeze severity weights before observing results; test plausible alternative weights and charge increased work with its actual sign.

[AI] Composite-endpoint research motivates component reporting because equally counted events can differ greatly in consequence. Here the implication is operational: a net-family metric guards against relabeling, while a mature disposition vector guards against overlooking a real recovery.

Basis: registered family rule and scope forks [HR]; conservation and rate identities [RC/AI]; Ferreira-González et al.2007 by methodological analogy. No historical intervention benefit is estimated.

TECHNICAL ANALYSIS

Matching totals do not identify which records share a status

[AI] Reason code and last status are different variables. Equal totals are a necessary reconciliation check when populations are intended to match; they neither verify entity membership nor supply a joint cross-tabulation.

Synthetic joint table	NULL + stalled	NULL + other	Known + stalled	Known + other
A: lower overlap	10	10	80	0
B: upper overlap	20	0	70	10

[AI/RC] Both hypothetical tables contain 100 units, 20 NULL reasons and 90 stalled statuses. Yet the stalled share of NULL records ranges from 50% to 100%. Let $a=|A|$, $b=|B|$ and N be the universe size. Then $\max(0, a+b-N) \leq |A \cap B| \leq \min(a,b)$, provided A and B are subsets of the same verified universe. This is an identification bound, not a confidence interval or a causal ownership share.

[HR/AI] The source compares reason-code and status totals, then interprets the unresolved reason segment. Even a tight overlap bound cannot establish current physical custody or the response that would have worked. The same literal in two taxonomies may describe different events; validate definitions before joining it.

[AI/RC] Missing stage attribution gives a second conditional bound. If c is linked coverage and s is the stage share among linked units, the full-population stage share lies between cs and $cs+(1-c)$. This requires correct links, aligned populations and valid classifications. Selective missingness can move the result anywhere within those bounds.

[AI] The registered coverage can conditionally preserve a majority recorded-location conclusion even under unfavorable unlinked records. It does not preserve the stronger assertion that those units remained there, were recoverable, or were preventable by the current owner.

[PD] Release the joint claim only after a unique-entity join, explicit unmatched states and a frozen as-of extract. Resolve duplicate weights and scope filters before narrowing a bound; never substitute equal row totals for that work.

Basis: exact synthetic extremal tables [AI/RC]; source marginal comparisons [HR]; partial-identification reasoning. Source counts, coverage coordinates and reconstruction keys remain private.

TECHNICAL ANALYSIS

Measure actionability independently of the label being challenged

[PD] Define the reference question before sampling: at the decision time, did this episode have a feasible, authorized response capable of changing its physical outcome? A reporting node, missing scan or accepted assignment cannot answer that question alone.

Construct	Measurement contract
History and custody	Separate last observed scan, actual location at decision time, event time, posting time and extraction vintage. Preserve unknown and disputed states.
Actionability and authority	Specify the available action, required resources, delivery window and owner’s decision rights. Responsibility to investigate is distinct from responsibility for causing the failure.
Outcome and reference	Adjudicate receipt independently at the original promise plus the fixed horizon. Retain nonreceipt, late receipt and on-time receipt; unresolved truth receives bounds, not automatic success.
Lineage and joins	Retain the entity key, filter, query, source version, multiplicity and reason for every exclusion. Query lineage explains a result’s origin without certifying physical truth.

[PD] Draw a random reference core across linked and unlinked records, reason/status combinations, business lines and node types. Supplement rare disagreements with known inclusion probabilities. Use design weights for population estimates; oversampled examples alone do not estimate network prevalence.

[PD] A hypothetical core of 300 independent histories yields a worst-case raw-agreement screen of about ±5.7 percentage points. Kappa precision depends on prevalence, coder margins, clustering and weighting; simulate it from a pilot. Blind coders to preferred ownership, adjudicate disagreements and report reference nonresponse separately.

[AI] A small known carrier-coded category does not bound all upstream causes unless the coding mechanism is sufficiently sensitive and complete. It is a descriptive comparison, not a valid causal positive-control falsification by itself. Give the favored local-control interpretation an explicit opportunity to fail.

Basis: proposed measurement design [PD]; Buneman et al.2001; Cohen1960; Lipsitch et al.2010. Control-design principles do not turn a convenient category into an independently validated control.

TECHNICAL ANALYSIS

Close the handoff and the correction, not just the analytical ticket

[HR/AI] The register describes repeated source-access attempts, divergent analytical outputs, explicit retractions and a raw export that overturned an absence claim. These support a disciplined reconciliation effort while exposing an incomplete delivery chain.

Acceptance stage [PD]	Evidence needed before credit
Evidence packet	The owner can retrieve the same eligible records and distinguish observed facts, inference and unresolved alternatives.
Accepted response	A role with authority acknowledges the action, scope, capacity and deadline; a named owner in an analytical document is insufficient.
Implemented correction	The repaired rule or document version replaces affected downstream copies, with confirmation from actual consumers.
Physical readback	The original cohort receives independent follow-up. Delivery, benefit and cash realization remain separate claims.

[HR/AI] An interface that omits a population may be capped, filtered, stale or incomplete. The registered independent retrieval changed the structural-absence conclusion. A logged correction or offered rebuild, however, does not establish that every shipped artifact was replaced or its readers updated.

[AI] Multiple agent reviews can share the same source, prompt and preferred interpretation. Agreement then supplies neither independent measurement nor independent causal confirmation. Require source-linked alternatives, disconfirming cases and an adjudicator who can retain an unresolved result.

[AI/PD] Handoff research highlights the work required to bridge information and responsibility across boundaries. Tucker–Edmondson adds a distinction between a workaround and removal of the recurring cause. AAC therefore requires both an accepted next action and verification that corrections reached the operating process.

[PD] Maintain a versioned action and correction log with consumer acknowledgement, unresolved dependencies and reopening triggers. Preserve current customer promises and refund entitlements. A separate, unrun refund-clock thesis is not authority to change those rules or to claim recovery savings here.

Basis: registered retrieval/retraction/handoff events [HR]; Cook et al.2000 and Tucker–Edmondson2003 by organizational analogy; proposed AAC acceptance tests [PD].

TECHNICAL ANALYSIS

Separate short-window movement from selection and maturation

[HR/AI] The source combines recent stage attribution, earlier last-status cuts and older taxonomy explanations. The youngest outcome period can have less time to accumulate records. Neither a large apparent movement nor a clean family total resolves these timing differences.

Variation and observation opportunity

[PD] Use compatible ordered observations only after defining the unit, exposure, extraction vintage and stable phase. Tentative XmR uses $MR_t = |x_t - x_{t-1}|$, $\text{mean}(x) \pm 2.66 \times \text{mean}(MR)$, and an MR upper limit of $3.267 \times \text{mean}(MR)$. These arithmetic limits do not identify causes.

[HR/RC/AI] The private audit contains 31 screens: 23 temporal screens, six same-period source/population forks and two artifact-evolution screens. The temporal group contains two four-week station-count snapshots, 18 two-period snapshots and three derived paired measures. No point or moving range crosses its tentative limits; no stable phase or eligible Laney panel is established.

[PD] Keep event occurrence, original promise, claim recognition and extraction dates separately. Compare equally mature cohorts or use explicit maturity models with validation. Do not infer the late-period count decline is improvement when claims may still arrive. Reconstructed prior values and repeated reports are not independent observations.

Selection and heterogeneous response

[HR/AI] The source prioritizes large bucket and site concentrations; no untreated repeat ranking establishes reliability or the counterfactual for the selected worst group. Under a hypothetical reliability of 0.30/0.60/0.90, expected untreated reversion is 70%/40%/10% of an initial excess. Those are sensitivity assumptions, not estimates for this case.

[PD] Freeze selected membership and retain an eligible comparison population. Model business line, node capability, route type, pre-existing backlog, source completeness and baseline volume as prespecified moderators. A volume concentration identifies workload, whereas a rate concentration identifies a different question.

[AI/PD] A shared recovery team creates interference: helping one domain can displace another. Block or enlarge the assignment boundary accordingly and measure queue effects. A regression adjustment cannot repair a violated capacity or outcome-ascertainment contract.

Basis: registered mixed-vintage and scope caveats [HR]; proposed SPC/selection plan [PD]; NIST; Laney; Barnett. No historical causal improvement, stable phase or regression-to-the-mean fraction is asserted.

TECHNICAL ANALYSIS

Attribute the analytical contribution at its documented stage

[HR/AI] The register supports a role in framing the decision, commissioning analysis, reconciling divergent claims and integrating a bounded action plan. It does not establish personal ownership of every query, operating remedy or downstream outcome.

Claim	Process-tracing grade and limit
Framing and integration [HR/AI]	Provisional registry hoop support for questions, synthesis, standards and decisions. Source authentication would strengthen provenance, not supply a customer counterfactual.
Numerical challenge [HR/AI]	Recorded re-derivation and retractions support analytical contribution. A desired controllability narrative still requires independent disconfirmation.
Delivery and adoption [HR/AI]	A created plan, named owner or persistent file does not establish accepted handoff, implemented repair or use in decisions.
Outcome and value [HR/AI]	No mature postintervention comparison identifies recovery effectiveness or realized savings. Category cost is context, not attributable opportunity.

Consolidated limits and reproducibility

[AI] Contribution rests on register transcriptions unless the underlying originals are recovered and authenticated. Equal totals do not establish common membership; last stage does not establish current custody; missingness and code sensitivity remain uncertain. Conflicting count, rate and scope branches are preserved rather than blended.

[AI] The available records do not identify actual route compliance, treatment versions, resource opportunity cost, mature causal response or cash release. A source’s claim that a fix is not being pursued cannot prove no work occurred outside the supplied corpus. Current completion status may remain unknown.

[PD] Freeze source hashes, query/filter logic, joint record crosswalks, weight definitions, as-of cuts, exclusions and alternative hypotheses. Link each public conclusion to its private calculation and claim class. Keep source identifiers, site ranks and numerical reconstruction anchors private.

[PD] A corrected document must carry its version to its consumers. A source-reproduction pass verifies the computation under its stated inputs; it does not promote causal interpretation or proposed implementation to historical fact.

Basis: selected role/governance records [HR]; Collier process-tracing discipline [AI]; proposed reproducibility contract [PD]. Original recovery status is maintained separately from analytical completeness.

TECHNICAL ANALYSIS

Test response routing at equal capacity on the original cohort

[PD] Randomize independent response domains to the evidence-and-owner-routing offer or approved usual response. The offer changes evidence and coordination, while both arms keep the same scheduled response slots and customer-entitlement rules.

[PD] The hypothetical design enrolls 200 objectively eligible, pre-outcome exception episodes per domain in eight cohorts of 25. Freeze membership upstream of routing. Include incorrect routes, no response, unresolved reference truth and all later outcomes. Adjudicate receipt at original promise plus 35 days; a response or reroute does not restart the clock.

[PD] Primary benefit is usual-minus-offer mean total variable burden per original episode: prespecified physical-outcome costs plus response and recovery effort. Charge fixed coordination separately. Nonreceipt reduction is a joint superiority gate; report late receipt and on-time receipt as components. Accepted routing is a mediator, not the analysis denominator.

[PD] The synthetic model assumes latent actionability and baseline disposition are independent conditional on domain type. Neither is inferred from a stage label. It enumerates possible episode outcomes and shared domain heterogeneity. Extra work can increase cost on some episodes, so improved mean value is not a claim of stochastic dominance for total burden.

[PD] Compare break-even and MDE in common reference units, using variance at the economic boundary and a conservative joint decision criterion. Gaussian planning remains an approximation to the finite process; simulate actual boundary laws, then pilot covariance, reference quality, actionability and resource constraints before recruitment.

[PD] Independently size damage, diversion, queue-delay and missing-reference safeguards. The working law does not power those gates. Unknown outcomes require worst-case or justified missingness bounds. Later-horizon receipts define a different estimand and require recalculation, not retrospective success coding.

[PD] If random rollout is unavailable, a separate staggered design can target group-time ATT against not-yet-treated domains under support, no anticipation and defensible interference assumptions. Event-study and equivalence diagnostics inform parallel trends; nonsignificance is not proof. Customer valuation needs mature field data and its own power design.

Basis: proposed domain-randomized response policy and synthetic finite mechanism [PD]; Hernán-VanderWeele2011; Callaway-Sant'Anna2021. No executed intervention or demonstrated financial gain is claimed.

TECHNICAL ANALYSIS

Make the attribution contract survive four disconfirming tests

[AI/PD] AAC connects lineage, bounded attribution, accepted authority, delivered response and mature outcome verification. It preserves useful location evidence while refusing to convert a label or a reporting improvement into preventable customer burden.

Proposition [PD]	Observable challenge and boundary [PD]
P1: Linked evidence improves actionable routing.	A blinded held-out audit finds no prespecified useful routing gain with adequate precision, or unresolved joins/clock errors remain material. Location agreement alone cannot pass.
P2: Accepted ownership improves response delivery.	Acknowledged assignments fail to increase timely authorized delivery under matched capacity, or gains require unmeasured queue displacement. Named ownership is insufficient.
P3: The package improves net physical value.	An upper net-benefit bound below zero contradicts value; a wide interval is inconclusive. Failure of nonreceipt or other harm gates blocks scale independently.
P4: Corrections reach the decisions they were meant to repair.	A controlled replay or consumer audit still finds the superseded claim used after the declared repair. A logged correction without readback does not pass.

[AI] Provenance theory asks which records and transformations produced an output. Partial identification separates what marginals and missingness allow from what they identify. Control-design research requires an explicit reason why a comparison can detect the hypothesized bias. Together they deny a shortcut from matching totals to causal ownership.

[AI] Compound-treatment theory requires specifying the actual versions of evidence, authority and response before transporting an effect. Handoff and organizational-learning research explain why information delivery and remedy delivery are separate mechanisms. These sources motivate tests; they do not supply historical success here.

[AI/PD] Deming cautions against reacting to unqualified variation. Campbell/Goodhart motivate tracking label migration and displaced work when a bucket becomes a target. OECD/JRC guidance makes composite weights and sensitivity explicit; Ittner–Larcker connects measurement to decisions rather than metric accumulation.

[PD] Preregister meaningful effects, cost boundaries and guardrails. Scale requires lower benefit bounds above the relevant thresholds plus valid measurement. Independent information value selects among evidence-purchase decisions; it must not be added to the same operating benefit being evaluated.

Basis: proposed AAC propositions [PD]; six fresh primary sources plus continuing measurement/learning methods [AI]. This is a falsifiable design framework, not an established historical causal result.

PROSPECTIVE MODEL / PHYSICAL RECEIPT AND OWNER RESPONSE

Test the response package on original episodes, not relabeled buckets

[PD] All quantities on these four pages are hypothetical. Randomize independent response domains to evidence-and-owner routing or the current approved workflow. Freeze exception eligibility before outcomes and retain every original episode.

Fixed design	Specified planning law
Population and capacity	200 pre-outcome exception episodes per domain, eight weekly cohorts of 25. Both arms have one approved response slot per episode. No post-routing exclusions or automatic refund hold.
Independent outcomes	Observe on-time receipt, late receipt and nonreceipt at original promise + 35 days. Independently adjudicate receipt and cause; missing evidence remains unresolved with bounds.
Conditional mechanism	Domain actionability is 0.40 or 0.70, equally likely. Given an actionable cause, correct routing rises from 0.45 to 0.75; approved-action success is 0.65. Location alone does not establish actionability.

[PD] Baseline on-time / late / nonreceived probabilities are 0.20 / 0.40 / 0.40, independent of actionability in this toy law. A successful appropriate action moves half of late receipts to on-time; it moves nonreceipts 40% to on-time and 60% to late receipt.

MEASURE THE OUTCOME VECTOR

[AI/PD] A label-only move has zero physical value. Verified nonreceipt-to-late-receipt recovery can reduce harm while the combined late-or-nonreceived count stays unchanged. Neither transition is established by aggregate family movements.

[PD] Unknown cause is not local ownership. Refund and customer-entitlement policy remain unchanged. The 35-day endpoint is proposed and requires maturity calibration.

PROSPECTIVE MODEL / CLUSTERING AND DECISION POWER

Require detectable break-even and a verified nonreceipt benefit

[PD] Exact finite episode and domain-mixture moments feed a conditional Gaussian planning approximation. A variance envelope is recomputed across routing efficacy; independent categorical simulation checks selected boundaries. These are not achieved field powers.

Design	Domains / arm	MDE (index)	Break-even detection	Joint power floor
Initial	16	7.17	26.0%	18.8%
Base candidate	80	3.17	86.1%	93.4%
Same / stress	80	3.95	67.7%	78.6%
Revised / stress	128	3.11	87.4%	93.6%

[PD] Usual variable resource burden is 100; fixed new cost is 3.44 index points. The two-sided 5% zero-effect test must detect that break-even at 80% power. Joint acceptance requires the approximate lower 97.5% bound on variable benefit to exceed cost and on nonreceipt reduction to exceed zero.

[PD] Domain actionability is stressed to 0.30 / 0.80 with the same mean. More episodes within one domain do not replace independent response domains. The joint floor is 1 minus the sum of the two marginal failure probabilities; no independence between outcomes is assumed.

CONDITIONAL RECRUITMENT, NOT A ROLLOUT COMMITMENT

[PD] The revised 256-domain design requires 51,200 original episodes and pilot-calibrated response mix, receipt truth and covariance. Recruitment is unestablished. Late-receipt, serious harm, refund compliance and workload safeguards need separate prespecified margins and sizing.

[PD] MDE and cost use index points, usual variable burden = 100. Gaussian bounds require field calibration and studentized/randomization inference; they are not exact finite confidence guarantees.

PROSPECTIVE MODEL / ECONOMICS AND MATURITY

Credit physical harm avoided and debit every additional response cost

[PD] Hypothetical harm weights are 0 / 4 / 18 for on-time / late / nonreceipt. Response effort is 0.9 for every attempt, plus 0.5 for an appropriate action and 0.6 for a recovered nonreceipt. These disjoint resource charges are not historical dollars.

[PD] Per original episode, avoided nonreceipt harm is 0.77220; added late-receipt harm is 0.01716 and added work 0.10824. Fixed cost 0.30 allocates coordination 0.12, evidence quality 0.10 and retry triage 0.08. Research-only costs are separate. $Net_raw = a \times h \times \text{avoided nonreceipt harm} - \text{added late harm} - \text{added work} - \text{fixed cost}$; $Net_index = 100 \times Net_raw / 8.7298$.

Avoidable share a	h = 0.80	h = 0.90	h = 1.00
0.60	-0.63	-0.10	0.43
0.80	0.79	1.50	2.20
1.00	2.20	3.09	3.97

[PD] The table reports net index points; a is avoidability and h the non-overlap share. Discount only the favorable nonreceipt component. Added late harm, work and fixed cost remain fully charged. If 25% or 50% of each arm’s day-35 nonreceipts arrive by longer follow-up, net becomes 0.19665 or 0.04650 raw units. These are different horizon estimands; reprice and re-power.

UNRESOLVED TRUTH CAN CHANGE THE DECISION

[PD] With 5% unresolved outcome evidence per arm and the observed-case contrast equal to this model’s contrast, worst-case nonreceipt bounds include zero. Do not delete unknowns. Field value scales the mature nonreceipt benefit and fully debits added late harm, response work and allocated cost on the same original volume.

[PD] Sensitivity shares and longer-horizon resolution are assumptions. Reconciliation and recovery benefits cannot be double-counted. Administrative relabeling alone earns no outcome value.

PROSPECTIVE MODEL / DECISION-INFORMATION VALUE

Buy evidence that changes routing, reconciliation or an investment hold

[PD] Decision-information value is a separate hypothetical utility model. The mutually exclusive state partition first assigns unresolved receipt/cause evidence; among resolved cases, distinguish an actionable physical cause from label-only correction.

Action	Actionable cause	Label only	Unresolved
Bounded routing trial	12	-8	-12
Reconcile metadata	1	8	-3
Acquire truth / hold investment	2	0	5

[PD] Priors are 0.45 / 0.30 / 0.25. The best prior action has utility 2.15; perfect information adds at most 6.90. Correct signals have probability q; each wrong signal has probability (1 - q)/2. Evaluate all 27 signal-to-action policies.

Signal quality q	EVSI	Evidence cost	Net EVSI
0.55	1.7475	0.40	1.3475
0.75	4.0375	1.20	2.8375
0.90	5.7550	2.40	3.3550

VALUE OF A DECISION IS NOT RECOVERED COST

[PD] These utilities describe choosing the right next investment, including reconciliation or an investment hold. They neither estimate historical decision value nor add to the operating-benefit model. Evidence acquisition cost is separate from subsequent action costs already included in the payoffs.

[PD] Priors, likelihoods and payoffs are declared sensitivities, not measured beliefs. Hard authorization, customer-entitlement and safety conditions remain mandatory.

Reproduce the population and ownership bridge before pricing a remedy

[AI] This is a diagnostic and action-planning case. It preserves the author’s analytical leadership while distinguishing the supplied historical record from new reconciliation, inference and proposed evaluation.

Claim class	Meaning and replication boundary
[HR] Historical record	An attributed statement or action in the selected project register. A prior statement that a query was independently recomputed is not the same as recovering and re-running that query in this case.
[RC] Recalculated	Exact checks on available source values, conditional missing-data and overlap bounds, process screens and wholly hypothetical prospective-model calculations. Matching arithmetic does not authenticate entity identity or field meaning.
[AI] Analytical inference	A conclusion tied to explicit assumptions: incomplete coverage, unidentified joint classifications, stale location, nonexhaustive codes, category migration or an unaccepted ownership handoff.
[PD] Proposed design	The Attribution-to-Action Contract, linked-record audit, accepted response package, prospective evaluation and economic/information models. These are not historical deployments.

Replication sequence

- [PD] Verify the source version and selected cells. Reconcile eligibility, node scope, entity grain, unit weights, event clocks, calendar windows and outcome maturity. Keep competing family and panel versions separate until their count and exposure bridges are recovered.
- [PD] Recompute matched-population bounds and feasible reason-by-status cross-tabs. Reconstruct every proposed treatment path, cost, mature outcome and sample-size gate independently. If the available data cannot identify a quantity, preserve the range or unresolved status rather than substituting a preferred point estimate.

THE DECISION THE CASE CHANGES

[AI/PD] Treat category names and stage observations as evidence for investigation. An accepted owner, verified action and later physical outcome are separate obligations. Only a measured response effect can support incremental operating value.

Public historical conclusions are qualitative and anonymized. Raw counts, denominators, source identities, queries, ranking detail and version-specific calculations are private.

RESEARCH REFERENCES / PROVENANCE AND OPERATING HANDOFFS

Preserve data lineage without confusing it with causal authority

Primary research and official methods inform this interpretation and proposed evaluation. They do not authenticate unavailable project sources, prove an owner assignment or supply an external effect size.

Buneman, P., Khanna, S., & Tan, W.-C. (2001). Why and Where: A Characterization of Data Provenance. ICDT, LNCS 1973, 316–330. DOI 10.1007/3-540-44503-X_20. Full author paper. Provenance helps trace how an output was derived. In this case, reproducible lineage and entity joins still do not prove that a label identifies a causal owner.

Cook, R. I., Render, M., & Woods, D. D. (2000). Gaps in the Continuity of Care and Progress on Patient Safety. *BMJ*, 320, 791–794. DOI 10.1136/bmj.320.7237.791. Indexed primary full text; direct access restricted. The operating analogy is a handoff requiring continuity, context and acceptance. No clinical outcome or benchmark is imported.

Lipsitch, M., Tchetgen Tchetgen, E., & Cohen, T. (2010). Negative Controls: A Tool for Detecting Confounding and Bias in Observational Studies. *Epidemiology*, 21, 383–388. DOI 10.1097/EDE.0b013e3181d61eeb. Indexed primary full text. A falsification check requires a justified relation to the bias it tests. A small named category cannot exclude unmeasured causes elsewhere.

Ferreira-González, I., et al. (2007). Problems With Use of Composite End Points in Cardiovascular Trials: Systematic Review of Randomised Controlled Trials. *BMJ*, 334, 786. DOI 10.1136/bmj.39136.682083.AE. Indexed primary paper. Different outcome components can have different importance and treatment responses. The case reports receipt states separately before valuing their proposed severity-weighted contrast.

Manski, C. F., & Molinari, F. (2021; online 2020). Estimating the COVID-19 Infection Rate: Anatomy of an Inference Problem. *Journal of Econometrics*, 220, 181–192. DOI 10.1016/j.jeconom.2020.04.041. Full author paper. Incomplete ascertainment motivates transparent partial identification. The logistics bounds here use their own coverage assumptions and do not transfer epidemiological rates.

Hernán, M. A., & VanderWeele, T. J. (2011). Compound Treatments and Transportability of Causal Inference. *Epidemiology*, 22, 368–377. DOI 10.1097/EDE.0b013e3182109296. Indexed primary full text. Different versions of an intervention can produce different effects. A routing label, evidence package and executed response therefore cannot be treated as the same treatment.

Primary-source access scope is recorded privately. Methodological analogies do not import clinical thresholds, benchmark effects or external certification into this operating case.

RESEARCH REFERENCES / MEASUREMENT AND PROSPECTIVE EVALUATION

Separate a process signal from a treatment and a decision threshold

Primary research and official methods inform this interpretation and proposed evaluation. They do not authenticate unavailable project sources, prove an owner assignment or supply an external effect size.

NIST/SEMATECH. e-Handbook of Statistical Methods, section 6.3.2.2: Individuals Control Charts. Official method. Tentative moving-range limits require ordered comparable observations. Short weekly series, different versions and cross-category summaries need distinct interpretations; none independently proves a remedy caused a movement.

Laney, D. B. (2002). Improved Control Charts for Attributes. *Quality Engineering*, 14(4), 531–537. Primary publisher abstract. A large denominator does not cure incompatible populations or extra-binomial variation. The case requires repeated eligible counts and stable definitions before estimating an attribute-process dispersion adjustment.

Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20, 37–46. Primary publisher record. Blinded coding should report confusion cells, raw agreement, kappa and uncertainty. Agreement on an administrative label is distinct from agreement on current custody, actionable cause or mature receipt.

Barnett, A. G., van der Pols, J. C., & Dobson, A. J. (2005). Regression to the Mean: What It Is and How to Deal with It. *International Journal of Epidemiology*, 34, 215–220. Primary abstract. A worst-mover overlay selects unusually adverse observations; subsequent recovery is not automatically a treatment effect. Preserve selected membership and a valid comparison cohort.

Hudgens, M. G., & Halloran, M. E. (2008). Toward Causal Inference With Interference. *Journal of the American Statistical Association*, 103(482), 832–842. Primary manuscript. Shared dispatch, staff and recovery capacity can transmit effects across units. The proposed assignment boundary must contain dependencies or estimate spillovers explicitly.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, 225, 200–230. Author manuscript. Group-time effects inform a later staggered-rollout alternative with explicit comparison populations, treatment timing and assumptions. Its identification and power do not follow from the hypothetical randomized model.

Strong, M., Oakley, J. E., Brennan, A., & Breeze, P. (2015). Estimating the Expected Value of Sample Information Using the Probabilistic Sensitivity Analysis Sample. *Medical Decision Making*, 35, 570–583. Author repository. Information is valuable through the decision it changes. The separate finite-state utility model is not a realized operating saving.

Primary-source access scope is recorded privately. Methodological analogies do not import clinical thresholds, benchmark effects or external certification into this operating case.

RESEARCH REFERENCES / OPERATING KNOWLEDGE AND CONTRIBUTION

Make measurement useful by connecting it to verifiable action

Primary research and official methods inform this interpretation and proposed evaluation. They do not authenticate unavailable project sources, prove an owner assignment or supply an external effect size.

Collier, D. (2011). Understanding Process Tracing. *PS: Political Science & Politics*, 44, 823–830. Primary publication. Contribution claims need discriminating evidence. Commissioning analysis, independently checking a result, integrating artifacts and implementing a remedy receive separate grades; derivative summaries do not become independent witnesses.

Tucker, A. L., & Edmondson, A. C. (2003). Why Hospitals Don't Learn from Failures. *California Management Review*, 45(2), 55–72. Author manuscript. Repeated local workarounds can leave underlying causes intact. The case uses this operating analogy to require remedy evidence and later outcomes rather than treating a closed ticket as learning.

The W. Edwards Deming Institute. The Funnel Experiment. Official exposition. Reacting to ordinary variation can increase process instability. The case separates reporting repairs, process signals and causal effects; an adverse category movement alone does not justify a remedy.

Campbell, D. T. (1976). Assessing the Impact of Planned Social Change. Original institutional record. Consequential indicators can change behavior and recording. Retaining the complete outcome family and independent receipt evidence reduces the temptation to treat a category transfer as improvement. No historical intent is inferred.

Goodhart, C. A. E. Problems of Monetary Management: The U.K. Experience. 1975 conference; revised volume listed as 1976 in the RBA archive. A measure's relation to its underlying construct can change when targeted. This case preserves physical customer states separately from labels, stage codes and ownership assignments.

Ittner, C. D., Larcker, D. F., & Randall, T. (2003). Performance Implications of Strategic Performance Measurement in Financial Services Firms. *Accounting, Organizations and Society*, 28, 715–741. Author manuscript. The strategic value of measurement depends on its operating use. A more detailed dashboard does not establish a better decision or a mature benefit.

OECD / European Commission JRC (2008). Handbook on Constructing Composite Indicators. Official handbook. Normalization, weighting, aggregation and sensitivity affect an indicator's meaning. The prospective severity score exposes component weights and reports physical outcome contrasts rather than allowing a composite to hide their differences.

Primary-source access scope is recorded privately. Methodological analogies do not import clinical thresholds, benchmark effects or external certification into this operating case.

FINAL REFEREE REPORT / INTERNAL AGENT-ASSISTED REVIEW

Accept the diagnostic contribution within its causal and delivery limits

[AI] This is author-led, agent-assisted review, not external peer review, client approval or a consulting-firm ranking. Acceptance concerns a bounded analytical case and proposed evaluation; it does not certify historical implementation or benefit.

Major issue	Disposition in the revised case
Location and coverage	Separated observed stage, current custody, actionable cause and control. Retained unmatched units and bounded whole-population stage shares under explicit alignment assumptions.
Marginals and source versions	Required an entity join before assigning missing reasons from another field. Preserved competing family/rate snapshots; a combined-country total remains a scope hypothesis to verify.
Observation and correction	Rejected structural absence inferred from an incomplete returned view. Distinguished a recorded correction from confirmed replacement of affected executive artifacts.
Evaluation and value	Defined a response package at equal capacity on original episodes. Retained mature receipt components and incremental effort; label-only transfers earn no benefit and existing entitlement timing remains unchanged.
Attribution and ownership	Credited commissioning, adjudication, integration and analytical restraint with provisional registry support. Preserved separate fix owners, unexecuted tests and unresolved delivery.

Partner critique: can a COO act on page one?

[AI] Yes: commission the linked diagnostic and an accepted ownership protocol. Data reconciles entities and clocks; Operations verifies the present exception and accepts the remedy; Quality and the evaluator measure later customer consequences.

Minor revisions and remaining limits

[AI] Preserve the preferred-hypothesis challenge, source-version ledger and corrected artifact trail. Calibrate treatment delivery, outcome truth, costs, dependence and recruitment before using the hypothetical power model for a funding decision.

REVIEW DISPOSITION / PASS WITHIN STATED SCOPE

[AI] The case supports analytical leadership, measurement reconciliation and a testable action contract. It establishes no settled cause for every category, shipped operating fix, realized saving or validated customer-policy change.

Scope: internal source, scientific, numerical and executive review. All claims remain bounded by the selected record and explicit hypothetical analysis.