

Customer Defect Analytics & Prioritization

Reconcile the defect tree, then validate the mechanism before choosing a fix

An anonymized operating case on reconciling fragmented defect categories, responding to executive review, and connecting an analytical finding to a testable action and accountable owner.

PROJECT CONTEXT / DEFECT DIAGNOSIS, PERFORMANCE REVIEW & OPERATING ACCOUNTABILITY

BLUF / BOTTOM LINE UP FRONT

Reconcile the full defect tree.
Validate the mechanism before assigning a fix.

ANALYTICAL FINDING

Offsetting category movements establish a net balance. They do not prove that the same orders changed state, locate the physical cause, or show that a proposed remedy worked. Each step requires different evidence.

CONTRIBUTION IN THE SOURCE RECORD

Directed cross-family analysis and the review response. Provisional contribution; operational acceptance and execution remain unverified.

EVIDENCE STATUS

Original artifacts were not recovered locally. The register supports a diagnostic deliverable; no operational or financial benefit is attributed.

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

STRUCTURED ABSTRACT / THE EXECUTIVE DECISION

Use the decomposition to focus the next test, not certify a cause

[AI] A disciplined review can replace a broad operating narrative with a specific, auditable question. Its value depends on keeping the accounting result separate from the mechanism, target population and authority to act.

Element	Decision-relevant finding
Context [HR]	An executive reviewer challenged a weekly defect paper to identify the actual growing subcategories and connect them to appropriate inputs. Repeated leaf labels appeared beneath several completion-state families.
Objective [HR]	Resolve the family-level ambiguity, answer the review comments and identify a narrower operating investigation with named owners. The deliverable was a diagnostic response, not a field experiment.
Approach [HR/RC]	Reconcile each canonical leaf across sibling families; compare volume with rate; retain competing versions; inspect candidate explanations and map each question to its evidence and proposed action.
Findings [RC/AI]	A conspicuous single-family movement largely offsets elsewhere. The retained arithmetic identifies a different growing group, but marginal offsets cannot identify parcel transitions and coded tags cannot authenticate physical cause.
Decision value [AI/PD]	Commission linked-cohort and station-level validation before targeting the proposed remedy. The analytical lead owns reconciliation; the operating owner accepts the test; the reviewer closes a question only against its stated evidence.
Implications [PD]	Require complete disjoint-leaf coverage, reconciled exposure, an explicit mechanism test and accepted operational ownership. Retain unresolved questions in the denominator. A closed comment does not certify a completed action.

ACTION FROM THIS CASE

[PD] Approve a bounded diagnostic test under the Reconciled Driver Decision Contract (RDDC). Set data, owner and decision thresholds before naming a target cohort; release causal or financial claims only after their own evaluation.

[HR] Source S1: selected project register; original artifacts remain unrecovered. [RC] arithmetic; [AI] interpretation; [PD] proposed design. Raw values and identity mappings are private.


EXHIBIT 1 / RESTORE THE FULL ACCOUNTING PERIMETER

Sum the same canonical reason across every mutually exclusive family

[HR/AI] The source’s useful methodological contribution was to leave the familiar family view and reconcile a repeated leaf across the complete hierarchy. That prevents a local decline from being mistaken for a network-wide reduction.

ONE REASON / FOUR EXCLUSIVE COMPLETION FAMILIES

Wholly hypothetical counts / constant meaning, complete coverage and a common cohort assumed



Reconciliation gate	Required evidence
Canonical meaning	Use a versioned reason dictionary, not a text-label join alone. Identical names can hide different definitions; changed names can refer to the same construct.
Exclusive records	Confirm that each eligible record belongs to exactly one sibling family and one relevant leaf. Repeated events, parent totals and children must not be counted together.
Complete perimeter	Retain unclassified and missing-family records, reconciled to the parent and the original eligible cohort. An unexplained residual stays visible.
Comparable versions	Freeze period, extraction vintage, maturity and eligibility before comparing totals. A later revision is not automatically the authoritative value for an earlier decision.

[AI] The source describes a reusable cross-family rule. Writing that rule is a contribution to analytical practice; it is not evidence that every later review followed it or that a downstream intervention succeeded.

FIRST DECISION

[PD] Accept a driver table only after its leaf totals reconcile to its parent. A partial export can generate a plausible ranking while omitting the category that would change the ranking.

[HR] Source S1: selected project register; original artifacts remain unrecovered. [RC] arithmetic; [AI] interpretation; [PD] proposed design. Raw values and identity mappings are private. Numerical diagram is wholly hypothetical.

EXHIBIT 2 / IDENTIFY WHAT MARGINS CANNOT REVEAL

Offsetting margins do not identify which orders changed state

[AI] The source interpreted opposing family movements as the same failure being booked to a different completion state. Even under the stronger assumption of a fixed closed cohort, margins alone leave multiple transition histories possible.

IDENTICAL MARGINS / DIFFERENT TRANSITION HISTORIES

Wholly hypothetical fixed cohort / group A: 80 → 60; group B: 20 → 40

WORLD 1			WORLD 2		
	To A	To B		To A	To B
From A	60	20	From A	40	40
From B	0	20	From B	20	0
Gross transitions = 20; net A → B = 20			Gross transitions = 60; net A → B = 20		

Basis: exact feasible transition tables. Across different weekly cohorts, linkage must be established first.

Same observed margins	Different unobserved histories
One-way movement [RC]	Of 100 synthetic records, group A moves from 80 to 60 and B from 20 to 40. Twenty A records move to B; none moves back. Gross movement is 20.
Two-way movement [RC]	The same margins arise when forty A records move to B and twenty B records move to A. Gross movement is 60. The net shift remains 20.
Actual weekly comparison [AI]	Successive weekly cohorts need not contain the same records. Without linked identifiers and a common risk set, even the fixed-cohort transition question is not yet defined.

[RC/AI] The sharp feasible range for A-to-B movement in this constructed cohort is 20–40. Netting solves an accounting problem; entity linkage solves a different identification problem. A near-canceling historical total cannot prove recoding, recovery or physical relocation.

WHAT CHANGES THE NEXT STEP

[PD] Request a linked transition matrix with entrants, exits and recodings. Keep the family reconciliation as a screening result while withholding any claim that lost orders were converted into late deliveries.

[HR] Source S1: selected project register; original artifacts remain unrecovered. [RC] arithmetic; [AI] interpretation; [PD] proposed design. Raw values and identity mappings are private. All counts and transitions in the exhibit are invented examples, independently enumerated.

EXHIBIT 3 / SEPARATE LOAD FROM INCIDENCE

More defect units can coexist with a lower defect rate

[HR/RC] The register reports a higher defect count with a slightly lower printed rate. The direction is robust to the recorded rounding branches, while the actual opportunity denominator is absent. An implied denominator is an algebraic check, not recovered exposure.

DEFECT VOLUME COMBINES EXPOSURE AND INCIDENCE

Wholly hypothetical indices / each quantity has its own baseline = 100

EXPOSURE
110

×

INCIDENCE
95 / 100

=

DEFECT VOLUME
104.5

Index identity: $110 \times 95 / 100 = 104.5$
A lower rate and a higher count can both be correct.

Basis: synthetic accounting identity. Historical implied exposures are not recovered denominators.

[RC] In the hypothetical example, exposure grows 10% and incidence falls 5%; defect volume therefore grows 4.5%. The identity is count = exposure × incidence. Using the prior rate first allocates +10 index points to load and −5.5 to incidence at the new load. A different allocation order reallocates the interaction without changing the total.

Interpretation	Decision implication
Customer burden	A lower rate does not erase the additional defect units. Count and rate answer different questions and require an explicit eligible population.
Process performance	Two rounded weekly rates cannot establish a sustained change or a treatment effect. Tentative process limits remain non-diagnostic with such a short sequence.
Subgroup inference	A total-concession segment split does not identify that segment’s late-delivery contribution. Require the exact outcome-by-segment cross-tab and common exposures.

DO NOT TRANSFER THE WRONG SEGMENT ANSWER

[AI] The source’s segment hypothesis can prioritize a data request. It cannot justify a segment-only remedy until the late-delivery outcome is measured at that segment’s grain.

[HR] Source S1: selected project register; original artifacts remain unrecovered. [RC] arithmetic; [AI] interpretation; [PD] proposed design. Raw values and identity mappings are private. The plotted indices are hypothetical; they are not transformations of internal source levels.

EXHIBIT 4 / DISTINGUISH DECOMPOSITION FROM ATTRIBUTION

A contribution above 100% reveals offsets, not a causal share

[HR/RC] The nominated growing group exceeds the net increase because other categories decline. That is valid signed component arithmetic when the categories are disjoint and exhaustive. It does not measure what would have happened without the proposed mechanism.

SIGNED COMPONENTS EXPLAIN A NET CHANGE

Wholly hypothetical change units / focal leaves +30, all other leaves -10, total +20



150% + (-50%) = 100%. Accounting closure does not identify an intervention effect.

Basis: synthetic disjoint components. Nested groupings cannot be added; residuals cannot be suppressed.

Accounting boundary	How to preserve it
Signed closure	In the invented example, +30 from the focal group and -10 from the rest produce +20 overall: 150% and -50% of the net increase. Retain both sides and any residual.
Nested grouping	A broader execution grouping already contains the focal timing leaves. Adding the two totals double counts those leaves; alternative views are not incremental opportunities.
Conflicting source sets	The source’s loss narrative and table use different component lists; one substitutes a late-delivery category. Keep both versions, but do not describe a mixed-outcome numerator as a valid loss share.
Counterfactual	A category can be net-flat while an intervention prevents a larger rise. Zero observed accounting contribution is not zero causal burden or zero preventable harm.

NOMINATE A TESTABLE MECHANISM

[PD] Treat the largest validated net grower as a diagnostic priority. Estimate the remedy’s effect using eligible treated and comparison units, mature outcomes and explicit concurrent-program controls.

[HR] Source S1: selected project register; original artifacts remain unrecovered. [RC] arithmetic; [AI] interpretation; [PD] proposed design. Raw values and identity mappings are private. All plotted quantities are synthetic. No historical decline is attributed to the diagnostic deliverable.


EXHIBIT 5 / TEST THE MEASUREMENT BEFORE CLOSING THE CAUSE

A falling reason tag cannot rule out an unmeasured causal path

[HR/AI] The paper checked several familiar explanations against their recorded tags. Those checks narrow the observed accounting story. Their causal force depends on sensitivity, specificity, timing and whether the nominated path actually maps to that tag.

THE SAME DETECTED SHARE CAN HIDE DIFFERENT CAUSE BURDENS

Wholly hypothetical / 100 records, 2 true-positive tags and no false positives



Source test	Bounded conclusion
Environmental tag declined	The recorded tag did not account for the net increase. Unmeasured or miscoded environmental effects remain possible; an adjacent address category is not automatically explained either way.
Input rose; output leaf fell	The two recorded series diverged. A cohort, maturity, taxonomy or mix mismatch can create that pattern; it neither establishes nor excludes a causal input effect.
Explicit execution tag was small	A label may omit failures recorded under other reasons. The analyst-created adjacent grouping needs validation before being treated as a physical mechanism.
Detection share was small	A small detected population is not an opportunity ceiling without verified coverage. Reporting location also differs from physical origin and operational addressability.

REOPEN ON MEASUREMENT EVIDENCE

[PD] Audit raw events and independent evidence, including records without the tag. Reopen the hypothesis when tag validity, population linkage or the causal mechanism fails its registered test; do not wait for the tag alone to rise.

[HR] Source S1: selected project register; original artifacts remain unrecovered. [RC] arithmetic; [AI] interpretation; [PD] proposed design. Raw values and identity mappings are private. Diagram assumes a synthetic population and perfect specificity only to expose sensitivity dependence.


EXHIBIT 6 / CONVERT THE ANSWER INTO AN EXECUTABLE DECISION

Name the evidence, owner and acceptance test for each review question

[HR] The response mapped reviewer questions to proposed actions and deadlines. That creates useful traceability. The register does not establish accepted station lists, completed field actions or verified follow-up outcomes.

FOUR DISTINCT COMPLETIONS / FOUR DISTINCT EVIDENCE TESTS

Proposed decision record / a written answer is an intermediate state



Decision record	Release condition
Analytical answer	Canonical leaves, signed reconciliation, common denominator, version fork and unresolved residual are independently checked.
Operating hypothesis	The physical path and controllable lever are stated; a reason-code name is not treated as the mechanism itself.
Target cohort	The operating owner approves the actual eligible station list and comparison design. A proposed worst-N cohort is not a recovered list or a causal control.
Action acceptance	The receiving owner explicitly accepts scope, resources and authority. The analyst cannot substitute a named owner line for acceptance.
Closure	Evidence distinguishes answer delivered, answer accepted, action completed and outcome validated. Reopened or unresolved questions remain in the original cohort.

[AI] A revision to the paper’s headline can improve precision without introducing new evidence. The retained body must use the same mechanism and cohort; a narrower heading does not repair an unsupported causal narrative below it.

[HR] Source S1: selected project register; original artifacts remain unrecovered. [RC] arithmetic; [AI] interpretation; [PD] proposed design. Raw values and identity mappings are private.

EXHIBIT 7 / PRESERVE CONTRIBUTION AND TEST THE NEXT DECISION

Credit the analytical work while testing the operational handoff

[HR/AI] The supported contribution is a stronger diagnostic and review response: a cross-family decomposition, review-question traceability, a reusable method and proposed ownership. This case separately audits source conflicts. Execution and mature customer effects require their own evidence.

RECONCILED DRIVER DECISION CONTRACT / PROPOSED

Define the question, test its evidence and route a decision under explicit authority



An unresolved cause can still have a useful next-evidence owner.
Evaluate decision quality before claiming downstream operating impact.

Basis: proposed RDDC workflow. Brand and methods are Luna Sol; historical client and system identities are withheld.

Contribution claim	Process-tracing grade and boundary
Reframed and directed analysis	Provisional registry-level hoop: the register describes the question, directed parallel analysis and resulting method. Same-session repetition is not independent corroboration.
Authored and revised response	Provisional registry-level hoop: described document metadata and a narrow version change support the account, pending recovery of the actual files and comments. Metadata alone cannot prove sole intellectual authorship.
Assigned follow-up ownership	Provisional registry-level hoop for the written proposals. Operating acceptance, the station lists and completion evidence remain unestablished.
Improved customer or cash outcomes	Unestablished. The diagnostic uses adjacent operating metrics; it does not identify an incremental mature effect or a realized economic benefit.

THE NEXT STUDY

[PD] Test whether RDDC improves the correctness and usefulness of driver-to-owner decisions while charging review, delay, rework and wrong action. Keep every original question; evaluate operational interventions separately.


[HR] Source S1: selected project register; original artifacts remain unrecovered. [RC] arithmetic; [AI] interpretation; [PD] proposed design. Raw values and identity mappings are private.

TECHNICAL ANALYSIS


A reconciled bridge is a starting point for a causal question

[AI/PD] The proposed Reconciled Driver Decision Contract (RDDC) separates a valid accounting bridge, a supported mechanism and an accepted operating action. Reviewer closure cannot substitute for any of these gates.

1 OPERATING EVENTS AND BOOKED ACCOUNTING



2 PROPOSED CONTROLLED DECISION BENCHMARK



[AI] Operating exposure X affects physical outcome Y and recorded reason G; outcome also affects completion state F. Coding practice and context C/U influence what reaches the booked bridge B. Summing B consistently does not invert this system to reveal X. A completion-state shift may reflect recovery, recoding, different cohorts or several mechanisms together.

[PD] The separate proposed benchmark randomizes reporting workflow Z. A visible triage signal S informs query Q; the hidden constructed state T affects its returned evidence and the reference J. This leads to an allocation or investigation decision R. An independently adjudicated reference J judges decision loss L; query effort W also costs resources. Fixed coordination K is charged once. The reference concerns evidence-supported decisions, not access to individual causal counterfactuals.

[PD] A later operating action requires a separate original cohort and mature comparison. The benchmark changes no customer route, refund or authority; decision-resource benefit is not customer savings.

Basis: registered decomposition and review workflow [HR]; proposed graph and benchmark [AI/PD]. Edges are hypotheses, not historically identified effects.

TECHNICAL ANALYSIS

Validate the accounting universe before ranking its movements

[HR/AI] The register describes a useful correction: aggregate the same reason across all sibling completion states. This prevents a partial-family change from being presented as the whole reason’s movement. The correction is conditional on the measurement contract.

Required identity	Why the bridge can fail
Meaning and grain	A shared label must denote the same event concept and unit weight; a similarly named process tag is not automatically the same defect.
Partition and period	Sibling states must be mutually exclusive and collectively cover the intended cohort at the same observation time. Preserve unknown and unclassified states.
Hierarchy and lineage	Add leaves once; do not add a parent to its children, or a Late-family row to a Lost-only numerator. Retain original IDs and versioned mappings.
Exposure and maturity	Use common eligible exposure and comparable extraction age. Count changes and rate changes answer different questions.

[AI] For reason j and completion state f , descriptive count movement is $\Delta C_j = \sum_f (C_{jf,1} - C_{jf,0})$. If reasons partition the same total, $\sum_j \Delta C_j = \Delta C$. A contribution $\Delta C_j / \Delta C$ can exceed 100% when other reasons improve. It is undefined when total change is zero and unstable near zero. Report signed contributions and offsets instead of constraining each contribution to 0–100%.

[HR/AI] The source’s alternative shared-upstream branches differ in category membership, not only arithmetic. One includes a row from another outcome tree. Such a numerator cannot be accepted as a competing valid share until the event universe is reconciled. Averaging the branches would conceal the problem.

[AI] For count C , exposure N and rate r , $C=N \times r$ and the exact symmetric bridge is $\Delta C = \Delta N \times (r_1 + r_0) / 2 + \Delta r \times (N_1 + N_0) / 2$. Subscripts 0 and 1 denote the compared periods. This separates volume and rate accounting without proving why either changed. Rounded rates and an absent exposure table support only an implied-denominator check, not a verified denominator or adjusted causal comparison.

Basis: historical register and preserved scope conflicts [HR]; independent arithmetic [AI]; Das Gupta’s standardization/decomposition distinction. Raw historical inputs remain private.

TECHNICAL ANALYSIS

Opposing family movements do not identify parcel transitions

[AI] Even perfectly reconciled marginal totals leave the movement of individual entities unidentified. The following constructed example assumes a stronger data structure than the historical record: the same 100 unique entities observed in two exhaustive states before and after.

Hypothetical joint movement	World 1	World 2
A stays A	60	40
A moves to B	20	40
B moves to A	0	20
B stays B	20	0
Before margins: both worlds	80A +20B	80A +20B
After margins: both worlds	60A +40B	60A +40B

[AI] Both worlds have the same margins: before(A,B)=(80,20), after(A,B)=(60,40). Forward movement is 20–40 and reverse movement 0–20; gross movement can be 20–60 while net movement remains 20. For a fixed universe N, with a entities in A before and b in B after, the intersection is bounded by $\max(0,a+b-N) \leq n(A \text{ before} \cap B \text{ after}) \leq \min(a,b)$. Linked IDs, not matching totals, identify the actual table.

[AI] The source compares reporting periods rather than supplying a closed entity-transition panel. Its matching declines and increases therefore do not establish that the same parcels were recovered, attempted or reclassified. The apparent Lost-to-Late conversion requires a separately linked original cohort with physical receipt and promise histories; it cannot be reconstructed by subtracting two weekly totals.

[AI] Nor does a flat reason imply zero causal effect. A second hypothetical example shows 100 defects before and after in two worlds. In one, follow-up would have had 120 defects without a remedy, which prevented 20. In the other, follow-up would have had 100 and the remedy prevented none. Both bridges are flat; their treatment effects differ. Zero observed accounting contribution is not zero avoidable burden.

[PD] Use offsetting family changes to trigger an entity-level transition query. If that query cannot be run, keep “candidate state redistribution” as a hypothesis and preserve the operating question. Do not close a maintenance action merely because its net bucket did not grow.

Basis: independent hypothetical exact enumeration, 73 checks [AI]; Freedman’s ecological-inference and bounds discussion. No example is fitted to historical quantities.

TECHNICAL ANALYSIS

A falling tag does not falsify every pathway that could create it

[HR/AI] The historical close-out challenged available explanations using their named tags. That is useful descriptive discrimination, but a tag must have validated coverage and a justified causal relationship before its decline can rule out a mechanism.

[AI] For true binary-exposure prevalence p , the observed tag proportion is $q = \text{sensitivity} \times p + (1 - \text{specificity}) \times (1 - p)$. In an independent hypothetical example with perfect specificity, true cases rise $100 \rightarrow 120$ while sensitivity falls $0.8 \rightarrow 0.5$; recorded tags fall $80 \rightarrow 60$. This does not show that detection changed in the case. It proves that tag decline alone does not logically refute increased underlying exposure.

Recorded observation	Defensible implication
A cause-coded leaf falls	The measured leaf did not explain the positive net accounting change under the verified taxonomy. Uncoded or indirect effects remain possible.
An upstream input rises while its Late tag falls	Different clocks, mix, recovery, coding and exposure may separate input from landing. No causal exclusion follows without a tested bridge.
Total-concession segment rates diverge	Investigate segment mix and outcomes; those totals do not allocate a Late-specific change or provide a natural control.
Same-week correlation is weak	Report the estimate and uncertainty. Clock mismatch can attenuate an association but does not guarantee a null. Component arithmetic is not a causal replacement.

[AI] Completion states are potentially downstream of operations and customer outcome. Restricting or adjusting on “attempted,” “recovered” or a concession family can remove part of the total effect or induce selection. Rosenbaum and Hernán–Hernández-Díaz–Robins motivate retaining the original eligible population for total-effect claims and explicitly separating state-specific descriptions.

[PD] Predeclare what would discriminate each explanation: linked transition evidence, validated reason sensitivity, an exposure-normalized change, or a randomized response. Preserve rival mechanisms that remain consistent with the observations. A literal tag with almost no records is not a negative control merely because it is small.

Basis: source hypothesis tests [HR]; independently derived measurement identity [AI]; Rogan–Gladen, Rosenbaum and Hernán et al. applied within their stated scope.

TECHNICAL ANALYSIS

Audit semantic consistency and reviewer closure separately

[PD] The measurement audit should answer two questions: whether the bridge faithfully represents the eligible records, and whether a proposed decision is supported by evidence and accepted authority. A response document can pass the first editorial check while failing the second decision check.

Audit object	Required fields or test
Bridge record	Stable entity and event key; original promise; reason concept/version; exclusive state; unit weight; denominator membership; extraction vintage.
Transition claim	Same-entity linkage across clocks; reasons for entry, exit and missingness; independent physical receipt or unresolved disposition.
Decision docket	Question, admissible evidence, explored alternatives, classification uncertainty, proposed owner, acceptance and permitted next step.
Closure record	Answer delivered, reviewer accepted, evidence gap resolved, action accepted, implementation verified and mature result are distinct states.

[PD] Draw an independent random core of 200 action-context dockets from a future expanded frame, stratified by reason family, completion state, source availability and decision ambiguity. Supplement rare taxonomy exceptions for error discovery. Two blinded reviewers apply a frozen codebook; a third resolves disagreements without seeing the assigned workflow. Include “insufficient evidence” as a valid reference outcome.

[PD] Report raw agreement, a full confusion matrix, prevalence and unweighted kappa with intervals; cause category is nominal. Under independent sampling, 200 observations give a worst-case approximate 95% agreement half-width of 6.9 percentage points. This is a screening rationale, not kappa precision: shared sources/reviewers require design-effect adjustment, and pilot prevalence should drive kappa simulation. Do not pool oversampled exceptions into population estimates without design weights.

[AI] Agreement is not validity, particularly when reviewers share the same extract and assumptions. Test false confident attribution, missed legitimate abstention and owner-acceptance errors against an independent reference. Preserve an unresolved reference rather than forcing consensus. Repeated agent convergence is a computation check; it does not create new operating observations or authenticate missing originals.

Basis: proposed semantic and decision audit [PD]; Cohen agreement methods. Sampling counts are proposed, not historical. Independent reference quality remains a testable assumption.

TECHNICAL ANALYSIS

Separate sparse accounting comparisons from operating signals

[HR/AI] The record contains short comparisons and reported correlation summaries, not a verified post-intervention operating panel. Ordered observations, comparable maturity and stable definitions must be established before a process-control interpretation.

[PD] For each genuinely ordered compatible sequence, tentative XmR uses $MR_t = |x_t - x_{t-1}|$, individual limits $\text{mean}(x) \pm 2.66 \times \text{mean}(MR)$, and moving-range upper limit $3.267 \times \text{mean}(MR)$. A two-point pair has one moving range and cannot demonstrate stability. Cross-family categories, overlapping groups and alternative document versions are not an operating time axis.

[HR/AI] The private audit contains 34 mechanical two-point screens: 25 quoted two-period comparisons, one component-sum branch, two inferred-baseline sensitivities, five overlapping week/year-to-date comparisons and one document-version size pair. All arithmetic X/MR points fall within tentative limits; no stable phase or eligible Laney panel is established. Printed correlations do not reconstruct missing daily arrays; lagging changes the paired sample. No missing observations, effective sample size, intervention phase or mature follow-up may be invented. Until repeated binary counts and matching opportunities are validated, large displayed volumes alone do not justify a p or Laney p' chart.

[AI] The proposed worst-station lists are absent; no empirical reversion fraction is estimable. Under an assumed linear conditional-mean model, expected untreated follow-up is $m + \rho \times (x - m)$, where m is the population mean, x the selected baseline and ρ assumed persistence. Values 0.4/0.7/0.9 imply 60%/30%/10% reversion of selected excess. These are hypothetical sensitivities, not a universal identity or case estimates. Freeze selection separately and retain untreated comparators.

[PD] Prespecify moderators that could change the mechanism: product/service segment, node function, backlog, available recovery capacity, reason-code sensitivity and authority boundaries. A segment with a worse aggregate concession rate need not have worse Late-specific risk or a larger treatment effect. Standardize to a shared baseline mix and show uncertainty; a comparison segment is not automatically a causal control.

[AI] A small p-value, control-chart signal or reviewer agreement would still not identify the intervention effect. Conversely, a non-significant short-series correlation cannot certify no harm, no effect or an exact clock. The next test should reduce the relevant identification uncertainty, not simply add another version of the same bridge.

Basis: source reporting structure [HR]; NIST/Laney monitoring conditions, Barnett regression-to-the-mean caution and proposed validation [AI/PD]. Historical limits and inputs stay private.

TECHNICAL ANALYSIS

Credit the analytical discipline without certifying its causal story

[HR/AI] A provisional registry-level hoop supports contribution to question framing, reconciliation and review-response design. It does not authenticate individual execution from original files, demonstrate adoption of the standing rule, or establish that assigned actions improved operations.

Contribution claim	Evidence and boundary
Problem reframing	The register describes reviewer feedback that prompted leaf-level analysis. A comment-to-answer map supports documented responsiveness, not correctness of every answer.
Analytical correction	Cross-family netting and disclosure of blocked station data are useful safeguards. Transition and causal explanations still need independent tests.
Method documentation	A reusable rule and revised narrative are described. A saved instruction does not establish repeated use or reliable implementation.
Operating ownership	Named owners and exit dates are recorded; acceptance, target lists, execution and mature outcomes are unverified.

[AI] The source retains an important tension: a net-flat reason was described as something not to chase, while standard-work protection in that area was still assigned. These decisions can coexist if one concerns explaining this week’s change and the other concerns preventing a large baseline burden. The case should state those different objectives instead of calling the reason either fixed or irrelevant.

[HR/AI] Consolidated limits: original records and executable joins are not authenticated here; decomposition branches differ in membership; reason sensitivity and segment-specific outcomes are unverified; weekly populations do not establish entity transitions; completion states may be post-treatment; no historical control or mature implementation panel is supplied. Source claims of avoided misallocation remain plausible decision intent, not measured savings.

[PD] Close a reviewer comment when its answer and evidence boundary are accepted. Close a mechanism claim only when its discriminating test passes. Close an operating action only after accepted authority, verified execution and prespecified mature outcomes. Keep these statuses separately auditable so editorial completion does not silently become operational credit.

Basis: provisional source-register process tracing [HR/AI], with original authentication and outcome gaps retained. Collier informs the bounded contribution grade.

TECHNICAL ANALYSIS

Test evidence-to-owner decisions before claiming a customer effect

[PD] The proposed RDDC benchmark randomizes independent analyst/reviewer domains to usual top-net interpretation or a linked-evidence decision workflow. Both arms receive equal source access and authority; independent frozen batches come from the same eligibility frame.

[PD] Freeze all dockets before assignment, including those that ultimately support the initial owner, an alternative owner or only an unresolved cause with an accepted next-evidence owner. The hidden reference concerns what the constructed evidence supports. It is not a claim to know the unique physical cause of every historical event.

[PD] A common visible triage signal informs the one-query entitlement, available in both arms during intake days 0–7. Hidden truth cannot guide investigation. Returned packets may confirm, contradict or remain unresolved; an accepted next-evidence owner must be documented in that visible packet. Charge queries and returned work even when unhelpful. Preserve approved customer and operating routes during analytical review.

[PD] Primary ITT is usual-minus-RDDC mean decision-resource loss across every original docket at the original intake-day-14 deadline; blinded reference adjudication finishes by day 28. Charge wrong confident allocation, avoidable investigation, unresolved deadline burden and all query effort; fixed coordination is separate and counted once. Independently judge false confident routing over the original denominator. A correctly supported unknown with an accepted evidence owner can be a valid decision.

[PD] Cluster shared reviewers, source repairs and capacity to limit spillover. Blind adjudicators to assignment and prevent later reference information from entering an earlier decision. Prespecify missing and late submissions, treatment versions, hold costs and whether an owner actually accepts the action. The accompanying hypothetical model must size economic and false-routing gates jointly, compare MDE with break-even, and retain uncertainty in query quality and domain heterogeneity.

[PD] A later operational study could stagger accepted interventions across independent capacity domains and estimate group-time ATT versus not-yet-treated domains. Require no anticipation, common support, stable measurement, interference controls and joint pretrend/equivalence diagnostics. Freeze original manifests and mature customer outcomes. That separate field study remains unpowered; decision quality is not customer impact.

Basis: proposed controlled decision benchmark and separate unpowered field extension [PD]. All model probabilities, costs and allocations are hypothetical; no historic efficacy is assigned.

TECHNICAL ANALYSIS

Make reconciliation useful by making the decision rule falsifiable

[AI/PD] RDDC adds four gates to a useful accounting practice: semantic reconciliation, linked mechanism evidence, accepted authority and an outcome test. A polished bridge may pass the first gate while the correct decision remains to gather evidence.

Proposition	Discriminating observation and failure condition
P1: the bridge is invariant	Reclassification among exhaustive sibling states preserves the combined reason total; duplicate membership, concept drift or unmatched exposure invalidates the sum.
P2: netting improves decisions	On unseen adjudicated dockets, the complete bridge reduces wrong confident allocation versus a partial-family view; failure includes no material gain or new false closure.
P3: linkage earns mechanism language	Matched entity evidence supports the proposed transition and survives clock/selection checks; incompatible transitions or unresolved linkage defeat the claim.
P4: closure has operating content	Accepted ownership and verified action improve prespecified mature outcomes enough to cover cost; answered comments or lower query time alone do not satisfy this test.

[AI] Das Gupta separates composition/rate accounting; Freedman exposes missing information inside aggregate margins. Rogan–Gladen motivates validating the tag-to-construct bridge. Rosenbaum and Hernán et al. constrain completion-state adjustment. Nosek et al. distinguish a hypothesis generated by the current export from a prediction tested on new evidence. These methods qualify, rather than certify, the source’s driver nomination.

[AI] Deming discourages operational tampering with short-run reporting noise. Campbell and Goodhart make comment-closure counts an incentive risk: fast answers can crowd out unresolved evidence. Ittner–Larcker–Randall motivates measuring decision and operating value separately, while Tucker–Edmondson supports retaining contradictions as learning inputs rather than repairing only the narrative.

[PD] Use source, reference and authority requirements as noncompensatory gates, consistent with the OECD warning about composite construction. EVSI may prioritize a source pull, semantic audit or owner test, but it cannot authorize a prohibited action or be added to the same avoided loss twice. Preregister meaningful effects and reopen conditions; a failed test constrains the nominated mechanism, not every possible explanation.

Basis: proposed RDDC propositions [PD] and bounded synthesis of primary and continuing methods [AI]. No proposition is claimed historically confirmed.

Test evidence-resolved decisions, not comment counts

[PD] Randomize independent analyst–reviewer domains to the usual response or RDDC. Each receives an independent frozen batch of 18 constructed dockets under the same information and authority contract. Every original docket remains, including returned, unanswered and unresolved work.

[PD] Both arms have equal information rights and one-query entitlement. Query in days 0–7; lock the final response on day 14; finish blinded adjudication by day 28. Returned evidence documents any accepted next-evidence owner. Reference states concern evidence sufficiency, not individual counterfactual causation. Current customer and response routes continue unchanged.

Decision mechanism	Wholly hypothetical law
Visible triage	Both arms see H/M/L. Usual queries H; RDDC additionally queries M. Hidden reference type is never a policy input.
Linked-ID query	One query costs 2 resource units and may confirm, redirect, or return unresolved. Return/rework penalties are included once in terminal loss.
Decision consequences	Wrong confident allocations cost 12 or 15; unowned holds cost 6 or 7. Correctly owned uncertainty costs 3; it is not false closure.
Unit of inference	Shared domain difficulty correlates dockets. Independent domains, not repeated questions, determine recruitment.

[PD] Base mean burden falls from 7.469 to 4.351 resource units per docket before the separate 0.75 protocol allocation. Unsupported allocations fall from 50.55% to 14.45%. These are invented planning outcomes. An all-hold policy costs 6.30 before setup and cannot match the modeled protocol.

DECISION STANDARD

[PD] Decision: calibrate reference validity, query returns and domain covariance before recruiting a confirmatory benchmark.

[PD] Synthetic constructed-case ITT. No actual owner reassignment, customer intervention, adoption or project savings is established.

A valuable response can still require an impractical test

[PD] Use independent-arm finite domain laws. The response path adds M-signal queries probabilistically; it is not a compliance estimate. Extra queries can harm already-supported cases, so stochastic dominance is not assumed.

[PD] For n domains per arm, exact rational moments bound the contrast variance over the entire response family. Bernstein radii provide a conservative lower confidence bound. The sufficient 80% MDE equals the 97.5% confidence radius plus the 80% detection radius; compare it with break-even, 10.04 index points.

[PD] Radius $r\alpha = RL/(3n) + \sqrt{[(RL/3n)^2 + 2VL/n]}$, $L = \log(1/\alpha)$. V bounds summed arm variances; R = 17 for loss and 1 for unsupported allocation.

Design	Per arm	MDE index	BE power floor	Joint floor
Initial	48	37.55	0.0%	47.3%
Base candidate	480	9.70	84.9%	>99.9%
Same size / stress	480	15.34	2.5%	>99.9%
Conditional redesign	1120	9.84	83.0%	>99.9%

[PD] Passing requires the MDE to be no larger than break-even and at least 80% joint power for positive net value and fewer unsupported allocations. Joint power uses a union bound. The two-sided zero test has at most 5% error; economic and wrong-allocation lower bounds each use 97.5% confidence within the specified bounded law.

[PD] Exact finite-polynomial diagnostics give 98.6% break-even detection for the base candidate and 59.4% under stress. They check this invented law; they do not validate field nuisance parameters.

DECISION STANDARD

[PD] The redesign requires 2,240 independent domains and 40,320 original docket. Recruitment remains unestablished; this is a mathematical allocation, not an operational recommendation.

[PD] Usual decision burden = 100. Conservative minima are 451 and 1078 domains/arm. All inputs require pilot calibration.

Discount avoided error burden; charge every added query

[PD] Value is decision-resource burden per original docket. Favorable terminal-loss reductions total 4.131 resource units; over-investigation adds 0.063 terminal loss and 0.95 query expense. Fixed K = 0.75 covers setup, owner coordination and recurring assurance over the same service pool.

[PD] $Net_raw = a \times h \times 4.131 - 0.063 - 0.95 - K$. $Net_index = 100 \times Net_raw / 7.469$. Here a is avoidability and h is non-overlap. Added harm and costs are fully charged. Returned work is already included in terminal penalties.

Avoidability a	h = 0.60	h = 0.80	h = 1.00
0.50	-7.01	-1.48	4.05
0.75	1.28	9.58	17.88
1.00	9.58	20.64	31.70

[PD] Entries are net index points at K = 0.75. With a = h = 1, K = 0.50 / 0.75 / 1.25 yields 35.05 / 31.70 / 25.01 points. Reprice and re-power lower-realization scenarios.

[PD] Ideal reference truth is assumed. If at most ε of each arm is unknown or mislabeled, loss-contrast uncertainty is bounded by 34ε raw units and unsupported-allocation uncertainty by 2ε. At ε = 5%, the expected net lower bound is 0.668 raw units. Missing truth is not zero loss.

[PD] Future field net = eligible treated volume × mature customer ATT improvement × avoidable cost/outcome × non-overlap – incremental cost. These inputs are unavailable; no customer valuation or field power follows.

DECISION STANDARD

[PD] Preserve the approved current response while evidence is unresolved. Neither a query nor a written answer authorizes causal blame.

[PD] Hypothetical resource units; usual burden = 100. The benchmark is randomized ITT; the separate observational field proposal targets ATT.

PROPOSED INFORMATION VALUE / 4 OF 4

Buy evidence only when it changes the authorized decision

[PD] A separate decision model compares a bounded trial, evidence repair and deferring investment while retaining the approved response. The mutually exclusive states concern evidence sufficiency, repairable linkage or unresolved scope/authority. No positive benchmark effect is assumed.

Action / state	Sufficient	Repairable	Unresolved
Bounded decision trial	14	-12	-18
Repair evidence linkage	4	11	-7
Retain route / defer	1	3	8

[PD] Priors are 0.40 / 0.35 / 0.25. A three-category diagnostic reports the true state with accuracy q and each other state with probability (1 – q)/2. Utilities include subsequent action costs; acquisition expense is separate.

Signal accuracy	Gross EVSI	Acquisition cost	Net EVSI
0.50	0.8750	0.50	0.3750
0.75	3.9625	1.25	2.7125
0.90	6.2350	2.50	3.7350

[PD] Enumerating all 27 signal-to-action policies gives prior value 3.7000 and EVPI 7.7500. Gross EVSI never exceeds EVPI. These are arbitrary decision-utility units, separate from the benchmark burden and never added to its savings.

DECISION STANDARD

[PD] Evidence may justify repair or an investment hold. It does not override mandatory evidence-admissibility and accepted-owner gates.

[PD] Low/base/high signal scenarios are assumptions, not estimated audit accuracy. Information value is independent of claimed project impact.

Reproduce the accounting result before testing its operational interpretation

[AI] The Reconciled Driver Decision Contract (RDDC) connects a complete accounting perimeter, an explicit mechanism hypothesis, accepted ownership and a prespecified decision test. Each link has a different evidence requirement.

Claim class	Meaning and replication boundary
[HR] Historical record	The supplied register describes a reviewed paper, a cross-family decomposition and a response with proposed owners. Its quoted documents and metadata remain derivative until the originals are recovered.
[RC] Recalculated	Source arithmetic, competing versions and tentative process screens; independently constructed transition bounds, count-rate identities and measurement counterexamples. Synthetic examples are not historical estimates.
[AI] Analytical inference	The full leaf sum improves the accounting answer. Margins do not identify parcel transitions; recorded causes do not identify physical mechanisms; named owners and due dates do not prove operational acceptance or execution.
[PD] Proposed design	RDDC and a controlled driver-to-owner decision benchmark. All original questions remain, including returned or unresolved answers. Wrong routing, unnecessary queries, delay and rework carry costs.

Replication sequence

[PD] Freeze the workbook, source pointers and selected cells. Recover the reviewed paper, comments, competing response versions and export; compare text and metadata before grading authorship. Reconcile canonical leaves, missing families and parent totals at a fixed vintage and cohort.

[PD] Retain every credible rate-rounding and component-set branch, including a branch that fails population admissibility. Validate reason tags and linked transitions independently. Obtain the actual station cohort and accepted owner before moving from a diagnostic to an intervention.

THE INTERPRETATION THIS CASE CHANGES

[AI] A more complete decomposition can overturn a family-level story while leaving the causal story unresolved. The method is strongest when it produces a narrower question and a better test, rather than declaring the mechanism solved.

Raw values, source identities, internal references and confidential version details remain in the separate private annex. Public numerical constructions and planning assumptions are hypothetical.

RESEARCH REFERENCES / AGGREGATION AND IDENTIFICATION

Do not infer individual mechanisms from aggregate margins

Primary methods research informs taxonomy reconciliation, causal boundaries and the proposed decision workflow. It does not authenticate historical originals or establish downstream operating impact.

Das Gupta, P. (1993). Standardization and Decomposition of Rates: A User's Manual. U.S. Bureau of the Census, Current Population Reports P23-186. Original official manual and overview. Separate additive volume and rate accounting from causal explanation. Exposure weights and the interaction allocation must be stated; an implied denominator is not an observed one.

Freedman, D. A. (1999). Ecological Inference and the Ecological Fallacy. University of California, Berkeley, Department of Statistics, Technical Report 549. Original manuscript and institutional record. Group margins leave individual transitions underidentified. The exact transition tables in this case demonstrate the ambiguity independently; no historical linkage is inferred.

Rogan, W. J., & Gladen, B. (1978). Estimating Prevalence from the Results of a Screening Test. *American Journal of Epidemiology*, 107(1), 71-76. DOI 10.1093/oxfordjournals.aje.a112510. Publisher metadata and original abstract. Observed positive tags differ from underlying cause prevalence when detection is imperfect. The case does not assume a diagnostic correction is identified without sensitivity and specificity.

Rosenbaum, P. R. (1984). The Consequences of Adjustment for a Concomitant Variable That Has Been Affected by the Treatment. *Journal of the Royal Statistical Society, Series A (General)*, 147(5), 656-666. DOI 10.2307/2981697. Publisher summary; original scan located. Completion states affected by operations cannot automatically serve as baseline adjustment variables. Preserve the original cohort when defining total effects.

Hernán, M. A., Hernández-Díaz, S., & Robins, J. M. (2004). A Structural Approach to Selection Bias. *Epidemiology*, 15(5), 615-625. DOI 10.1097/01.ede.0000135174.63482.43. Original abstract and indexed first-page text. Conditioning on a common effect can induce selection. Booking or completion membership needs explicit causal treatment; no universal direction of bias is claimed.

Nosek, B. A., Ebersole, C. R., DeHaven, A. C., & Mellor, D. T. (2018). The Preregistration Revolution. *Proceedings of the National Academy of Sciences*, 115(11), 2600-2606. DOI 10.1073/pnas.1708274114. Full original article. A mechanism selected after inspecting one export remains exploratory. Prespecify the next discriminating test and keep discovery separate from confirmation.

Access scope is documented in the private research memo. Applications are methodological; no outside empirical performance is attributed to this project.

RESEARCH REFERENCES / EVALUATION AND UNCERTAINTY

Account for selection, shared reviewers and decision costs

Primary methods research informs taxonomy reconciliation, causal boundaries and the proposed decision workflow. It does not authenticate historical originals or establish downstream operating impact.

NIST/SEMATECH. e-Handbook of Statistical Methods, section 6.3.2.2: Individuals Control Charts. Official method. Two-period leaf comparisons supply arithmetic checks but no stable process phase. Computed tentative limits cannot certify a sustained improvement or its cause.

Laney, D. B. (2002). Improved Control Charts for Attributes. *Quality Engineering*, 14(4), 531–537. Primary publisher abstract. Attribute charts require eligible numerator-exposure pairs. A large network volume does not repair missing denominators, repeated records or changing taxonomy.

Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20, 37–46. Primary publisher record. Use independent blinded reason and owner coding, disagreement matrices and uncertainty. Agreement does not establish physical cause or accepted operating authority.

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-N station cohort selected from noisy performance can improve through reversion. Freeze selection, estimate reliability when data permit and retain nonselected comparisons.

Hudgens, M. G., & Halloran, M. E. (2008). Toward Causal Inference With Interference. *Journal of the American Statistical Association*, 103, 832–842. Primary manuscript. Shared analysts, reviewers and operating capacity create interference. Assign independent review domains and audit cross-domain learning instead of counting every docket as independent.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, 225, 200–230. Author manuscript. A separate future staggered operating rollout needs group-time effects, defensible comparisons and explicit anticipation and spillover assumptions. The decision-workflow benchmark does not power that field study.

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. Value new linkage or taxonomy evidence through the decisions it can change after acquisition cost. Keep this evidence-purchase utility separate from review-resource effects and customer benefit.

Duchi, J. (2026). *Statistics and Information Theory*. Author lecture notes, July 23, section 4.1.2, pp. 78-79, Proposition 4.1.21 and equation (4.1.10). Original statement and derivation verified. The independent bounded-summand Bernstein tail supplies the conservative confidence radius. The domain variance envelope is derived within this case; the notes do not establish recruitment feasibility or field calibration.

Access scope is documented in the private research memo. Applications are methodological; no outside empirical performance is attributed to this project.

RESEARCH REFERENCES / MEASUREMENT AND GOVERNANCE

Distinguish a completed answer from a useful operating decision

Primary methods research informs taxonomy reconciliation, causal boundaries and the proposed decision workflow. It does not authenticate historical originals or establish downstream operating impact.

Collier, D. (2011). Understanding Process Tracing. *PS: Political Science & Politics*, 44, 823–830. Primary publication. Distinguish framing, authorship, proposal, acceptance and implementation. Described metadata and self-reported method records provide provisional support pending recovered originals.

Tucker, A. L., & Edmondson, A. C. (2003). Why Hospitals Don't Learn from Failures. *California Management Review*, 45(2), 55–72. Author manuscript. Retain unresolved contradictions and returned review questions as opportunities for organizational learning. Closing a comment locally does not demonstrate a system-level operating correction.

The W. Edwards Deming Institute. The Funnel Experiment. Official exposition. Reacting to routine variation can worsen a process. A short-period defect ranking needs stability and mechanism checks before it becomes a program of repeated operational adjustments.

Campbell, D. T. (1976). Assessing the Impact of Planned Social Change. Original institutional record. A target for closed comments can change reporting behavior. Preserve unresolved and reopened questions so output volume cannot substitute for evidence quality.

Goodhart, C. A. E. Problems of Monetary Management: The U.K. Experience. 1975 conference; revised volume listed as 1976 in the RBA archive. A review score may lose its meaning when it becomes the target. Correct routing, justified uncertainty and accepted authority must remain visible alongside completion counts.

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. Link analytical measures to accountable decisions and test their value. A more detailed driver table is not evidence of operating improvement or financial return.

OECD / European Commission JRC (2008). Handbook on Constructing Composite Indicators. Official handbook. Do not allow a weighted review score to hide invalid source scope, unaccepted ownership or harmful action. Keep mandatory evidence and authority gates separate from aggregate utility.

Access scope is documented in the private research memo. Applications are methodological; no outside empirical performance is attributed to this project.

FINAL REFEREE REPORT / INTERNAL AGENT-ASSISTED REVIEW

Accept the diagnostic contribution; evaluate the remedy separately

[AI] This author-led, agent-assisted review is not external peer review or client approval. Acceptance concerns reproducible accounting, bounded inference and a proposed decision-quality test.

Major issue	Disposition in the revised case
Margins treated as transitions	Replaced the asserted parcel migration with sharp synthetic transition bounds. Across actual weekly cohorts, linkage and a common risk set remain unestablished.
Accounting treated as cause	Preserved the complete cross-family decomposition while separating net contribution from causal burden. Positive and negative offsets, nested sets and unassigned residuals remain visible.
Incompatible measurement	Retained count-versus-rate distinctions, rounding branches and the mixed-outcome component table. Whole-concession segment rates cannot allocate the late-delivery result; coded tags require independent validity checks.
Comment closure treated as execution	Separated answer delivery, reviewer acceptance, accepted operational ownership, completed action and mature outcome. Proposed station lists and instrument changes are not upgraded into verified adoption.
Prospective decision value	Specified a sequential query-and-owner benchmark with retained original questions and priced errors, investigation and unresolved work. Customer effects remain a separate future field estimand.

Partner critique: can a COO act on page one?

[AI] Yes: require complete reconciliation, commission a linked mechanism test and name the accepting operating owner. The case identifies which evidence must change the decision before assigning a narrow intervention or declaring a hypothesis closed.

Minor issues and remaining conditions

[AI] Recover originals, resolve the source chronology and action inventory, validate the taxonomy and recruit independent review domains. Reference quality, query performance and decision-resource prices require calibration before a field-use claim.

REVIEW DISPOSITION / PASS WITHIN STATED SCOPE

[AI] The case supports a bounded diagnostic contribution and the proposed RDDC. It establishes neither historical causal attribution nor operational adoption, customer benefit or financial return.

Internal source, methods, numerical and executive review. Release acceptance is limited to the stated evidence and model conditions; no outside endorsement or ranking is implied.