

OPERATING CASE STUDY / CUSTOMER EXPERIENCE & LOGISTICS

Customer Recovery & Remedy Design

Test physical recovery without treating a delayed refund as a service improvement

An anonymized operating case on a proposed exception-recovery mechanism, the limits of last-scan evidence, and the decision rights needed before changing an automated customer remedy.

PROJECT CONTEXT / PHYSICAL RECOVERY, CUSTOMER REMEDIES & INVESTMENT GATES

BLUF / BOTTOM LINE UP FRONT

Prove incremental recovery first.
Keep the customer's remedy clock intact.

ANALYTICAL FINDING

An in-station scan does not establish recoverability at the refund instant. Even a successful recovery-and-hold bundle cannot identify what the hold contributed. Test the operational components separately.

CONTRIBUTION IN THE SOURCE RECORD

Framed the question, directed analysis and documented a gated proposal. Provisional process-tracing support; sole authorship and operational ownership are not established.

EVIDENCE STATUS

Originals unrecovered. Design only: no verified implementation, assigned owner, completed timing test or realized saving.

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

STRUCTURED ABSTRACT

Separate a recovery opportunity from authority to delay a remedy

[AI] The project’s transferable contribution is a proposed operating mechanism and an explicit evidence gate. Its source record does not demonstrate that withholding a refund was necessary, beneficial or authorized.

Abstract element	Decision narrative
Context [HR]	Two concession categories carried station-stage records. A proposed mechanism combined a filter repair, more frequent exception recovery and a temporary refund hold. It remained without an assigned owner or implemented pilot.
Objective [PD]	Identify whether better detection and earlier physical recovery improve outcomes for the original flagged customers while preserving their existing remedy timetable. Evaluate any later remedy change as a separate policy intervention.
Approach [RC/AI]	Reconcile attributed and full populations; distinguish recorded stage from current custody; separate component effects, customer outcomes and accounting events; examine pilot transportability and incomplete financial assumptions.
Findings [AI]	Missing attribution and blank mechanisms compound coverage loss. A daily work cadence does not establish that a later refund beats recovery. A bundle result cannot identify the refund hold’s incremental contribution.
Decision value [PD]	Assign an accountable sponsor to a bounded state-and-timing validation, then test recovery under the existing remedy clock. Retain false flags, unsuccessful attempts, costs and unresolved outcomes in the evaluation.
Implications [PD]	Operations owns physical work; data owners authenticate event linkage; customer-policy owners control remedy changes; Finance validates avoidable costs. The proposal author’s analytical role does not confer those authorities.

THE FIRST INVESTMENT DECISION

[PD] Fund evidence acquisition and a calibrated recovery-only test when the prerequisites are met. Do not translate gross concession exposure, a last-scan label or a borrowed site result into committed savings.

Basis: source register S1; new RREC analytical framework. Labels: [HR] historical record; [RC] recalculated; [AI] analytical inference; [PD] proposed design. No historical causal result is established.


EXHIBIT 1 / STATE AND TIMING

A scan identifies a record; it does not prove a recoverable package

[HR] The proposal treated the last recorded station-stage classification and no intervening terminal scan as the basis for a recovery gate. Its refund-time verification was not completed. [AI] The proposed query would improve linkage, but absence of a recorded exit still would not prove physical custody or feasible recovery.

FOUR DISTINCT CLAIMS / FOUR DIFFERENT EVIDENCE REQUIREMENTS

Proposed evidence ladder / each transition can fail even when the preceding record is correct



Evidence step [PD]	What it must establish
Recorded state	Identify the item, event meaning, event time, ingestion time and clock timezone. Preserve late corrections and unmatched records. A stage classification is not an observed package position.
Current physical presence	Audit the last-scan proxy against independent observation at the decision instant. Include stale scans, mislinked items, unseen movements and terminal events that arrived after the extract.
Feasible customer recovery	Verify item condition, safe handling, transport capacity, original promises and whether the customer still wants the delivery. Presence alone does not establish any of these conditions.
Incremental result	Compare recovery actions with the existing process on original eligible customers. Receipt, remedy, reorder, return and customer burden may coexist; they are not a single binary outcome.

A TIMELINE IS A MEASUREMENT REQUIREMENT

[AI] A once-daily sweep can occur before or after a refund. Relative cadence and an approximate delay do not identify the ordering for the same item. Link actual eligibility, work, receipt and remedy times.

Basis: S1 timing thesis and uncompleted verification [HR]; state-validation ladder and ordering critique [AI/PD]. The diagram shows required distinctions, not measured transition probabilities.

EXHIBIT 2 / POPULATION COVERAGE

Missing attribution and blank mechanisms compound the diagnostic gap

[HR] The author required stage shares to be reported against attributed volume, because one category had incomplete tier coverage. [RC] Reconciliation adds a second loss: an attributed record may still lack a usable mechanism code. The full matched population is the correct starting point for a coverage audit.

MECHANISM COVERAGE / MATCHED PARENT POPULATION = 100
Historical source-transcribed counts normalized to the same full category / no raw values disclosed



Basis: S1 matched counts. Unclassified does not mean physically absent, unrecoverable or invisible to every list.

[RC] Conditional on the source counts describing the same matched universe, index the category total to 100: attributed records represent about 74.5 and attributed nonblank mechanisms about 54.0. Roughly 46 index points are outside that classified view. This is coverage, not the proportion physically missing or impossible to recover.

[AI] A high station-stage share within observed records does not identify the same share among uncovered records. If c is coverage and s is the observed stage share, the full-population recorded-stage share lies between $c \times s$ and $c \times s + 1 - c$, provided the covered records form a correctly matched subset. Physical recoverability remains another unknown.

Reconciliation rule [PD]	Operational consequence
Match scope before division	Align period, geography, node type, category, unit grain and record version. Do not treat a different dashboard population as the missing remainder.
Keep uncovered and blank distinct	A missing tier row and a blank mechanism within a tier row require different repairs. Count both without adding overlapping populations.
Use alternative signals	A missing mechanism code need not prevent a stage-based recovery list. Validate actual list inclusion, physical lookup and action completion independently.

Basis: S1 matched category, attributed and blank-mechanism counts [HR]; parent population = 100 transformation [RC]. Raw keys and exact arithmetic remain private. This is a cross-sectional coverage bridge, not a time series.

EXHIBIT 3 / COMPONENT ATTRIBUTION

A successful bundle would still leave the refund hold unproven

[HR] The proposed gate combined a detection repair, a faster work cadence and a refund hold. Its recovery precedent concerned manual exception work, not the hold. [AI] These components can act through different paths, and the operational benefit need not require a change to the customer’s remedy.

IDENTICAL BUNDLE RESULT / DIFFERENT CONTRIBUTION FROM THE HOLD
Wholly hypothetical coherent potential outcomes / deliveries among 100 original cases in each world

WORLD	USUAL	RECOVERY ONLY	RECOVERY + HOLD
World A	20	50	50
World B	20	20	50

Added hold effect under enhanced recovery: 0 in A, 30 in B. Bundle result alone cannot decide.

Basis: exact two-world proof. Current recovery + hold delivers 20 in both worlds; no actual hold is authorized.

[AI] Observing better receipt under recovery plus a hold than under usual operations identifies only that bundle contrast under the study assumptions. It does not reveal what would happen with the same recovery work and the usual refund clock. A refund-time presence query also cannot supply that missing counterfactual.

Contrast [PD]	Question it can answer
Recovery-only versus usual	Do better detection and earlier work improve customer-accepted physical outcomes, after all lookup, handling and verification costs, with existing remedies preserved?
Recovery plus hold versus recovery-only	If separately authorized, does changing remedy timing add enough customer-valued benefit to justify delayed relief and other harms? Keep recovery resources comparable.
Joint interaction	Does the value of a remedy change depend on the recovery mechanism? A permitted factorial design can estimate component interactions; an unrestricted main-effect claim cannot ignore them.

START WITH THE COMPONENT THE EVIDENCE CAN SUPPORT
[PD] The prospective model in this case tests recovery before the existing deadline. It credits no withheld refund. Any later hold requires its own authority, population, customer safeguards and evaluation.

Basis: S1 three-component proposal and separate-site precedent [HR]; component identification argument [AI]; proposed contrasts [PD]. Exact two-world proof appears in the technical appendix.


EXHIBIT 4 / SCALING ASSUMPTIONS

A site-stage precedent cannot validate a whole-network savings rate

[HR] The entitlement applied a recovery rate from exception-rack work to a broader concession base. The source disagreed on whether the precedent involved one site or two, and supplied no corresponding control population or baseline. The cited result belongs to another team.

LIMIT THE EXTRAPOLATION TO THE POPULATION ACTUALLY TESTED

Wholly hypothetical / 100 baseline failures; 40 in the tested stage; stage failure reduction = 50%



Basis: transportability counterexample. Historical effects require comparability, controls and outcome linkage.

[AI] Transport requires more than applying the same percentage. The precedent concerns a particular stall stage; the new population contains other stages, missing mechanisms, different physical states and different customer needs. A stalled-unit reduction must also be linked to incremental receipt and avoidable resource cost.

[RC] In a wholly hypothetical example, the tested stage accounts for 40 of 100 baseline failures and its intervention reduces those failures by 50%. If other stages receive no benefit, the overall failure reduction is 20%, not 50%. Extending the stage rate to everyone overstates that constructed effect by 2.5 times. Neither percentage is a historical project estimate.

Unpriced assumption [AI]	Evidence needed before commitment
Eligibility at the decision instant	Authenticate current state and actionable capacity on the same original cohort. A gross historical concession amount does not establish reachable value.
Effect outside the precedent	Estimate effects by stage, node type, customer promise and prior recovery process. Preserve nonfocus units and measure displacement.
Conversion into net value	Separate physical service gain, refund accounting and avoidable cost. Include new work, continued entitlements, inventory movement, repeated contacts and overlap with existing programs.

Basis: S1 explicitly modeled entitlement, unresolved precedent scope and missing investment cost [HR]. Exhibit arithmetic is invented [RC]. No source scenario, projection or adjacent-program benefit is a realized result of this project.

EXHIBIT 5 / DETECTION RELIABILITY

A blank mechanism code is not automatically an invisible package

[HR] The record identified both a blank-mechanism population and an alleged filter escape affecting a particular scan state. It treated those as recovery blind spots. [AI] Their combined effect depends on population overlap and the actual list-generation rule; adding the percentages would double-count or mix denominators.

DO NOT ADD A WHOLE-POPULATION SHARE TO A CONDITIONAL ESCAPE RATE

Wholly hypothetical / 100 flags; 30 blank codes; 20 in a scan class; 25% of that class escape = 5

LIST EXCLUDES BOTH GROUPS

Omitted union = 30 to 35
Captured = 65 to 70

VALID STAGE SIGNAL ADMITS BLANKS

Only 5 filter escapees omitted
Captured = 95

The overlap and the implemented list rule determine coverage.
Neither captured count is a delivered-package result.

Basis: set-union bounds with overlap from 0 to 5. Alternative-signal performance must be measured.

[RC] In the hypothetical exhibit, 30 of 100 flagged items have blank mechanism codes. Twenty belong to a scan-state class, of which five escape a filter. If the list excludes both blank-code items and escapees, between 30 and 35 are omitted, depending on overlap. If a validated stage signal admits blank-code items, the filter omission alone is five.

[AI] These calculations describe possible list membership, not physical recovery. Repeated alerts, duplicated packages, stale state, unreachable items and incomplete execution can further separate a generated list from actual service. Conversely, a blank field may be harmless if another reliable signal supplies the needed action.

Validation measure [PD]	Denominator that preserves meaning
List sensitivity and precision	Independently adjudicated actionable items and all listed items, respectively; include a sample of unlisted cases to measure false negatives.
Action completion	All assigned eligible items, with attempt time, outcome and reason for noncompletion. A touch or a scan is not successful receipt.
Risk prioritization	Show severity, occurrence and detection judgments separately. An ordinal FMEA product is a screening score, not a failure probability or financial expectation.

Basis: S1 detection hypotheses and estimated FMEA [HR]; invented overlap counterexample [RC]; proposed measurement contract [PD]. Historical filter escape was not independently reproduced as a completed repair.

EXHIBIT 6 / OUTCOME AND PROCESS CONTROL

A lower concession score can coexist with unchanged customer service

[HR] The proposed control plan tracked concession defects, list adherence and downstream category shifts. The gate never entered production. [AI] Changing remedy policy can directly alter the financial metric, so a chart shift cannot by itself establish recovered delivery or a positive customer effect.

PHYSICAL RECEIPT AND CUSTOMER REMEDY CAN COEXIST

Proposed linked ledger / classify both outcomes on the same original customer and observation horizon

RECEIPT + REMEDY
Possible entitled outcome

RECEIPT / NO REMEDY
Check policy and timing

NO RECEIPT + REMEDY
Relief without recovery

NEITHER VERIFIED
Unresolved customer need

Track original promises, remedy timing, accepted receipt, repeat work and mature cost.
A reduced refund count alone receives no physical-service credit.

Basis: RREC outcome-vector design. These categories do not establish entitlement rules or historical proportions.

Outcome family [PD]	What the evaluation must retain
Customer and physical service	Original promise, customer-accepted receipt, time to useful resolution, nonreceipt, reorder and cancellation. Report both event incidence and waiting-time burden at a fixed maturity horizon.
Customer remedy	Entitlement, actual issue time, reversals and unresolved requests. Receiving the item and receiving compensation are not necessarily mutually exclusive.
Operations and economics	Every search, failed attempt, handling touch, extra transport, verification task and displaced unit. Separate implementation costs from research-only measurement.
Data and program integrity	Original eligible denominator, unlinked records, late-arriving events, mechanism recoding and cross-program overlap. Keep false flags and unsuccessful cases in the assigned group.

MONITORING AND CAUSAL EVALUATION HAVE DIFFERENT JOBS

[PD] Use tentative process limits only for comparable, ordered observations. A planned target is not an observed point; a source-mentioned baseline without its values cannot be charted. Use randomization for the proposed treatment contrast.

Basis: S1 proposed control design and absence of implementation [HR]; outcome separation [AI/PD]. No historical post-treatment chart, attributable reduction or realized saving is asserted.


EXHIBIT 7 / OWNERSHIP AND RELEASE DECISIONS

An evidence gate needs an owner, a deadline and a permissible next action

[HR] The record specifies a verification condition before an executive savings commitment and carries the proposal into a formal design. It establishes no completed verification, assigned operating owner, pilot or sponsor endorsement. [AI] The stated gate constrained the proposed claim but supplied no mechanism for resolving it.

TURN AN OPEN ANALYTICAL GATE INTO AN ACCOUNTABLE DECISION

Proposed sequence / each release requires its own evidence and accountable owner



Decision [PD]	Accountable role and release evidence
Validate the opportunity	An assigned analytical owner authenticates source versions, state at the decision instant, missing records and feasible recovery. A sponsor sets a bounded evidence budget and review date.
Test recovery-only	Operations owns safe execution; the customer-policy owner confirms unchanged remedies; data and Finance owners approve measurement and cost definitions. Size the study against economic break-even.
Consider remedy changes	A separately authorized policy owner evaluates the incremental hold contrast and customer safeguards. The recovery-only result does not authorize or validate a hold.
Scale, revise or stop	An accountable decision forum reviews mature all-original-cohort results, costs, adverse effects and transportability. An unresolved gate leads to an explicit disposition, not silent approval.

PERSONAL CONTRIBUTION / PROVISIONAL PROCESS-TRACING SUPPORT

[HR/AI] The register supports question framing, directed analysis and proposal development, including a self-imposed claims gate. Workflow provenance limits sole-authorship inference. Neither named sponsors nor document production establishes operational approval or delivery.

Basis: S1 recorded proposal and unclosed dependencies [HR]; prospective ownership contract [PD]. Contribution is provisionally supported by derivative records; no smoking-gun or doubly decisive proof of implementation is claimed.


TECHNICAL ANALYSIS

A physical recovery and a customer remedy are different events


[AI/PD] The Recovery–Remedy Evidence Contract (RREC) tests an earlier recovery offer while preserving the existing remedy. A historical scan, a safe recovery opportunity and an accepted receipt require separate evidence.

RECOVERY OFFER / REMEDY POLICY / PHYSICAL REFERENCE

Proposed RREC graph. All arrows are hypotheses; no historical coefficient is estimated.



SEPARATE MEASUREMENT CHANNEL / A FLAG IS NOT PHYSICAL PROOF



P is part of C/U above. S defines the frozen flag cohort; independent V checks it before action.

False flags remain in the ITT. H may affect willingness; no refund hold or withheld-cash value is tested.

[PD] Z assigns the recovery-only offer; V verifies current eligibility; A records attempts; Y is timely, later or absent accepted recovery. Actual state P and capture Q generate the scan signal S. C/U includes capacity, condition, customer willingness and unmeasured causes. False flags remain in the assigned population.

[PD] Y and variable resources W determine burden B, measured through independent ascertainment J as O. Fixed setup cost K is charged outside B once. H is remedy policy and timing; it can affect customer experience and willingness, but Z does not change H in the proposed trial. Refund receipt F is a distinct event.

[AI] The graph separates measurement from intervention. No recorded terminal event is not proof of presence, and payment is not an absorbing physical state. Arrows express mechanisms to test; the register supplies no estimated intervention coefficient. The separate policy panel prevents recovery evidence from silently authorizing a hold.

Basis: registered recovery and remedy proposals [HR]; explicit proposed causal structure [AI/PD]. No historical implementation or causal effect is established.

TECHNICAL ANALYSIS

Reconstruct the risk set at the time a decision was possible

[HR/AI] The linchpin timing query is specified but unrun. A useful operational list may start with stage signals; estimating incremental recovery requires identity, state and clocks joined at the original decision.

Quantity [AI/PD]	Definition and identification boundary
Recorded pre-remedy flag	An episode satisfying the frozen signal rule using only information available before its existing remedy deadline.
Qualified recovery share	Independently present, safe, accessible and customer-accepted opportunities / all original flags. Presence alone is insufficient.
Timely accepted-recovery ITT	Offer-minus-usual accepted receipts before the unchanged remedy deadline / all original assigned flags.
Burden ITT	Usual-minus-offer mature service and variable-resource burden per original flag; subtract fixed incremental cost once.

[PD] Maintain event time, database-availability time, original promise, flag creation, physical inspection, attempt, customer acceptance, actual receipt and remedy execution. Deduplicate at an explicit package–episode grain; preserve authorized identity changes and multiple financial events separately. Use later truth to audit the decision, not to fabricate what was known beforehand.

[AI] Snodgrass and Ahn distinguish represented history from when the database held that history. Here, an old station scan can remain the latest known scan after an unrecorded departure. An as-of join tests visibility and timing; it still needs independent physical reference to establish current presence.

[AI] Daily cadence does not prove that refunds win the race. With the same sweep schedule, early visibility can permit a timely attempt while late visibility cannot. Compare each episode’s first actionable visibility and actual attempt with its own remedy deadline. Do not add overlapping loading and response intervals twice.

[PD] Align eligibility, assignment and follow-up before response. Excluding episodes found absent, already remedied or unsuccessfully searched after assignment would change the estimand. Time-to-attempt and list adherence are process diagnostics, not substitutes for accepted receipt.

Basis: unexecuted registered query [HR]; temporal database distinctions (Snodgrass–Ahn) and target-trial alignment (Hernán et al.) [PD]. Historical timestamps remain unauthenticated.

TECHNICAL ANALYSIS

A bundle result cannot identify the value of delaying a remedy

[RC/AI] Two coherent hypothetical worlds produce the same improvement for recovery plus a hold, yet disagree completely about whether the hold adds value. The missing component comparison matters.

Recoveries among 100 original cases	World A	World B
Current recovery; current remedy	20	20
Enhanced recovery; current remedy	50	20
Current recovery; remedy hold	20	20
Enhanced recovery; remedy hold	50	50
Added hold effect with enhanced recovery	0	30

[RC] Construct each world with 20 cases always recovered and 50 never recovered. The remaining 30 respond to enhanced recovery alone in A; they respond only when both components are present in B. These individual potential outcomes generate every table cell. Both worlds satisfy individual non-harm in this example, so the unidentified component is not an artifact of hidden offsetting harm.

[AI] Observing only current practice and the combined package cannot distinguish the worlds. Neither list adherence nor a lower recorded concession count supplies the missing counterfactual. An untested hold could also change customer willingness, contacts and waiting burden, which the simple proof deliberately leaves outside its recovery count.

[PD] The immediate comparison is enhanced recovery versus current recovery with the existing remedy unchanged. Any later hold test requires separate authority, customer-welfare margins and a design that identifies the relevant component contrast and interaction. No hold arm is authorized or powered here.

[AI] Collins and colleagues explain why package trials and factorial experiments answer different questions. Lack of a continuous response-surface model does not prohibit randomized comparisons. Conversely, a factorial design is not automatically the most feasible next step when one component lacks authority and a safe operating specification.

Basis: independently constructed exact potential-outcome proof [RC], reproduced privately; component-design interpretation [AI/PD]. All table inputs are invented and have no historical calibration.

TECHNICAL ANALYSIS

Coverage compounds before a mechanism can be tested

[HR/AI] A mechanism share among attributed records is not full-category coverage. Missing attribution and blank codes are separate gaps; neither necessarily makes a package invisible to every recovery signal.

[AI] Let c be the fraction linked to the attribution table, s the recorded station-stage share among linked records, and b the blank-mechanism fraction among linked records. Under a common universe and valid stage classification, whole-category stage share lies between $c \times s$ and $c \times s + (1 - c)$. Full-category nonblank mechanism coverage is $c \times (1 - b)$. A unit-based coverage bound cannot price dollar exposure without stratum-specific amounts.

[RC/HR] The registered counts imply that only just over half of the matched full category has both attribution and a nonblank mechanism code. This conditional reconciliation is materially weaker than quoting the nonblank share inside the attributed subset. It does not establish that the remaining records cannot be surfaced by another stage or exception signal.

Independent audit task [PD]	Reference and consequence
Physical presence at decision	Independent location evidence; distinguish an absent item from an absent observation.
Safe, accepted recovery opportunity	Condition, accessibility and current customer preference; unresolved cases remain unresolved.
Clock and identity linkage	Package—episode, as-of knowledge and remedy deadline; flag false and duplicate joins.
Receipt and remedy outcome	Independently confirm accepted receipt and actual payment timing; do not infer either from a work-list closure.

[PD] Start with an independent random core of 200 episode-context records plus labeled exception oversamples. Under independent sampling, worst-case raw agreement has an approximate 95% margin of 6.9 percentage points; shared site, source or package clusters require a design-effect revision. Kappa precision needs prevalence-specific pilot simulation, alongside confusion matrices and raw agreement.

[PD] Blind adjudicators to assignment and reconcile disagreements against independent evidence. High agreement can reproduce the same defective feed. Weight supplemental cases only under known sampling probabilities; bound unverified truth rather than treating it as absence or silently dropping it.

Basis: conditional coverage algebra [RC]; measurement taxonomy and proposed agreement study [AI/PD]. Cohen’s agreement coefficient does not establish truth or causal recoverability.

TECHNICAL ANALYSIS

A recovered queue item is not automatically a recovered customer

[HR/AI] The cited adjacent pilot concerns a particular backlog stage. It does not identify this project’s whole-category receipt effect, the contribution of a refund hold or a finance-approved saving.

[AI] Let w be the tested stage’s share of original baseline failures and r its relative failure reduction. A relative reduction of $w \times r$ in all baseline failures follows if the endpoints match, the stage effect is causal and other stages are unchanged. A share of all shipments would require baseline failure risks as well. Migration between categories and adverse workload effects need separate terms; a smaller list can mean departure, disposal, reclassification or receipt.

[AI] “Still in the building” is not a cash-conversion state. A customer may have purchased a replacement, declined the recovered item or retained a remedy entitlement. Delivery can occur after a refund. Payment suppression, physical recovery and service improvement therefore require separate outcome ledgers and decision rights.

Ledger [PD]	Value that may be supported
Physical and customer outcome	Timely accepted recovery, later accepted recovery, no verified recovery; original cohort and maturity preserved.
Operating resources	Every inspection, failed attempt, handling touch and capacity displacement, including false flags.
Customer remedy	Eligibility, actual execution time, contacts and waiting harm under the authorized policy.
Investment	Incremental fixed readiness and recurring costs; research-only evidence cost tracked separately.

[PD] Price the signed mature outcome vector against usual practice, discount favorable avoidable benefits where justified, and charge added work and harms in full. Fixed costs enter once. Do not add overlapping broad and attributed exposure pools, convert a gross concession balance into controllable value, or annualize a modeled percentage as realized savings.

[AI] Tax and colleagues distinguish complaint outcomes from procedure and treatment. Anderson and colleagues show that failure source changes customer evaluation in their setting. These motivate independent customer checks here; neither transports an effect size or establishes that a later item compensates for a delayed remedy.

Basis: source pilot and exposure limitations [HR/AI]; separate service, cost and remedy estimands [PD]. The proposed numerical study assigns no refund-cash benefit.

TECHNICAL ANALYSIS

Do not turn sparse descriptions into control or rollout evidence

[RC/AI] The private audit computes 24 mechanical screens. All points and moving ranges fall within their tentative arithmetic limits; no validated stable operating phase or intervention effect follows.

[RC] The inventory comprises three four-week count descriptions; two alternative-final-vintage sensitivities; three same-period scope comparisons; four source-version comparisons; one reported policy endpoint pair; two policy-restatement pairs; one mismatched-clock outcome pair; four adjacent-pilot maturity or definition pairs; three engineered target paths; and one sibling forecast-versus-manifestation pair.

[PD] For a comparable ordered scalar sequence, the provisional individuals limits use $\text{mean} \pm 2.66 \times \text{mean adjacent moving range}$; the moving-range upper limit is $3.267 \times$ that mean. The counts lack a validated repeated exposure and binary-unit process, so no operational Laney panel is eligible. A plan trajectory, mixed vintage or two-point comparison cannot establish routine variation.

[AI] The source's diffuse site shares also do not prove that only a network-wide rollout could matter. A maximum site share does not determine a selected group's combined share or prove that group insufficient. The upper bound $\min(1, k \times \text{maximum share})$ for k sites can be uninformative. Baseline share is not causal preventability, achievable recovery or staffing capacity.

[PD] Freeze any worst-performing-site selection on a separate baseline and report all selected sites, not only improvers. Repeated noisy rates invite regression to the mean. No empirical RTM fraction is identified here; illustrate reliability-dependent shrinkage as sensitivity, then validate on an independent period with matched seasonality and exposure.

[PD] Prespecify moderators: presence precision, age of scan, time remaining, package condition, business line, customer willingness and shared capacity. Use pretreatment values and interaction uncertainty. Selecting only inspected, accepted or successfully recovered items would condition on intervention mechanisms and lose the all-flag effect.

[AI] Deming's warning against reacting to noise applies only after a valid measurement boundary. A favorable plan or source-version movement cannot earn a process signal by being plotted. Monitoring flags investigation; randomized comparisons identify the proposed offer's effect.

Basis: private source-specific SPC inventory and calculations [RC]; no historical raw coordinates disclosed. Selection, RTM and moderator proposals are prospective [PD].

TECHNICAL ANALYSIS

Credit the design contribution without upgrading its evidence

[HR/AI] The register supports a recovery hypothesis, a gating query specification and a proposed control plan. It does not establish a completed query, accepted owner, deployed gate, mature outcome or approved remedy change.

[HR/AI] Commissioning, analytical direction and insistence on validation are plausible contributions documented through the register. Claims of sole authorship require original provenance because the recorded workflow also includes AI-assisted drafting and an unresolved go-ahead. A named sponsor, signature or self-review score does not prove endorsement, execution or independent assurance.

[AI] Process tracing therefore supports only a provisional hoop: the transcribed chronology is consistent with hypothesis formation and design influence. Original artifacts were not recovered in the bounded search. A stronger claim needs authenticated versions, instruction-to-change traces, execution records, ownership acceptance and independent outcome evidence. Absence from the supplied corpus does not prove that no work occurred elsewhere.

[AI] The intellectual connection is specific. Jeffery and colleagues' sensor-stream work exposes the tradeoff between missed observations and stale presence inference; it is an analogy, not a claim about the installed technology. Snodgrass and Ahn separate event history from database knowledge. Together they make an as-of list auditable without confusing it with physical truth.

[AI] Campbell and Goodhart explain why targeted counts can change meaning under pressure. Recovery-list closure, hold adoption and lower concessions can improve while accepted receipt or customer welfare does not. Composite scoring cannot let high adherence compensate for absent authority or unreliable physical reference. The measurement-performance association discussed by Ittner and colleagues is not a causal validation of this program.

[AI] Consolidated limits are material: uncertain source authentication; unmatched exposure and maturity branches; missing unit-level temporal linkage; untested state precision; nontransportable adjacent pilots; unspecified hold duration and owner; unmeasured labor and displacement; and no mature intervention panel. These limits bound both funding confidence and personal attribution.

[PD] Preserve the useful design discipline by converting each unknown into an explicit evidence gate. A negative result should be able to retire the proposed mechanism, narrow the target population or keep recovery separate from policy. Repeated agreement among agents or copied records is not independent corroboration.

Basis: derivative source chronology [HR]; bounded contribution analysis and substantive literature connections [AI]. No original source or implementation has been authenticated.

TECHNICAL ANALYSIS

Test a recovery-only offer on the original flagged population

[PD] Randomize matched independent recovery domains before any new inspection. The offer adds an earlier verified opportunity; usual practice retains the final sweep. Both keep the existing customer-remedy policy and deadline.

Protocol element	Prospective specification
Population and clock	All original pre-remedy flags; four cohorts, fixed exposure per domain. Observe accepted physical receipt through original promise plus 14 days.
Assignment and contamination	Randomize within matched domain pairs; independent staff and capacity across arms. Shared grades are a planning assumption, not an observed matching result.
Primary acceptance	All-flag service/resource burden improvement clears fixed cost, and timely accepted recovery improves; one endpoint cannot compensate for the other.
Missing or failed response	Keep false flags, unsafe contents, refusal, failed searches and missing disposition; bound unknown truth with all denominators fixed.

[PD] The numerical model uses wholly hypothetical presence, qualification and two-opportunity recovery probabilities. Its finite probability tree supplies exact moments; known-normal variance-envelope power remains an approximation checked by finite-process simulations. Lower presence and imperfect matching can defeat a design that passes the base case. More flags in one shared domain cannot replace independent domains.

[PD] A blinded calibration pilot must estimate physical-reference accuracy, current acceptance, between-domain covariance, attempt cost and recruitment feasibility. Match on pretreatment variables. Do not interpret the model’s response-law mixture as observed compliance or the causal effect of execution among successful cases. Research adjudication is common to both arms; production verification is charged once.

[PD] Freeze safety, customer-wait, refusal and ordinary-flow guardrails with their own authorized margins and power. The two modeled gates do not certify these outcomes. Prespecify symmetric urgent-harm reporting; routine research findings should not create differential extra recovery work before readout.

[AI] Neither a positive synthetic calculation nor a source backlog precedent authorizes rollout. If independent domains, valid truth or acceptable customer conditions cannot be obtained, redesign and re-power or leave the investment decision unresolved. A refund-hold trial remains a separate, unapproved question.

Basis: proposed RREC randomized ITT; full numerical assumptions and economic sensitivity follow [PD]. No historical probabilities, investment benefit or customer-policy authorization is asserted.

TECHNICAL ANALYSIS

Four propositions make the Recovery–Remedy Contract falsifiable

[PD] RREC is an evidence contract for a bounded next decision. It permits the favored mechanism to fail instead of treating a promising exposure estimate as a result.

Proposition	Discriminating test and failure condition
P1 / State fidelity	An as-of flag identifies current qualified recovery opportunity. Blind physical and customer reference must meet preregistered precision and missingness margins; otherwise repair the signal.
P2 / Timing opportunity	Earlier actionable visibility produces additional pre-remedy accepted receipt. A randomized recovery-only contrast must exceed a material threshold; cadence adherence alone cannot pass.
P3 / Net customer value	All-flag mature service benefit exceeds full incremental cost without violating independently powered harm gates. Payment deferral is not credited as benefit.
P4 / Component authority	A remedy change adds permitted value beyond recovery alone. Without a separately authorized identifying comparison and customer margins, leave this proposition unresolved.

[PD] Register endpoints, minimum relevant effects, missing-truth bounds and stop rules before viewing treatment contrasts. A nonsignificant estimate does not falsify a mechanism unless its confidence interval excludes the prespecified material effect. Conversely, a point estimate above break-even does not clear a lower-bound funding rule.

[PD] If randomization later becomes infeasible, a staggered rollout could estimate group-time ATT among treated domains against not-yet-treated domains. Require no anticipation, common support, aligned mature cohorts and an interference boundary. Joint event-study diagnostics and prespecified pretrend equivalence checks assess plausibility; they do not prove parallel trends. This extension has no power claim or customer-benefit estimate.

[PD] Evidence purchase asks which permitted decision would change after better state, clock or authority information. EVSI uses separate state utilities and acquisition costs; it is not added to service benefits. Refund-hold authority is a hard gate, not a utility term that a favorable average can waive.

[AI] The immediate decision is to authenticate the risk set and test verified recovery under the existing remedy clock. Retain the source’s concern about missed recovery opportunities while withdrawing its unsupported claims that recorded station state proves convertible loss or that recovery success establishes the need for a hold.

Basis: original RREC propositions and prospective decision rules [PD]; component and temporal identification limits [AI]. A future observational ATT differs from the proposed randomized ITT.

Test recovery under the existing remedy clock first

[PD] Randomize the whole recovery-list and earlier-sweep offer within matched independent operating-domain pairs. The comparator retains the current final sweep. Existing refund eligibility and timing stay unchanged; withheld payments generate no credited value.

Frozen element	Wholly hypothetical specification
Population	40 original pre-remedy flags per domain across four cohorts. Keep false flags, refusal, failed attempts and missing disposition. Follow each original promise through day 14.
Physical eligibility	Presence is 0.40 or 0.80 by shared pair grade; 0.75 of present items are safe, accessible and customer-accepted for recovery. These are separate latent conditions.
Recovery clock	A final sweep succeeds with probability 0.45 on qualified items. The offer adds an earlier 0.45 opportunity, then the same final sweep if unrecovered. Remaining qualified items have 0.30 later accepted recovery. Timely means before the existing remedy deadline, not the original promise.
Resource accounting	Each applicable inspection costs 0.18 and each accepted recovery 0.80 hypothetical resource units. Timely / later / absent recovery carries service burden 0 / 4 / 12. Added fixed cost is 0.40 per original flag.

[PD] Independent reference checks distinguish current presence, safe action, customer acceptance and actual receipt. Perfect reference classification is only a planning condition. No last scan is treated as proof of recoverability.

RECOVERY ASSIGNMENT IS NOT A REFUND-HOLD TEST

[PD] The mature all-flag intention-to-treat contrast measures service/resource burden and timely accepted recovery. A separate hold proposal would require customer-policy authority, waiting-harm margins, independent reference and its own powered comparison.

[PD] Equal fixed flag exposure makes equally weighted domain means match the per-flag estimand. No actual pilot, customer-policy change or historical transition probability is asserted.

A design that clears the base case can fail lower presence

[PD] Exact finite outcome moments retain the shared grade and repeated requests. Quadratic variance envelopes over the stated response family support approximate normal power; finite-process simulations check the approximation. They do not establish field calibration.

Design	Pairs	MDE index	BE detection	Joint floor	Decision
Initial	24	6.81	41.9%	69.4%	Fails
Base candidate	64	4.17	81.8%	98.1%	Conditional
Lower presence	64	3.67	90.2%	50.2%	Fails
Imperfect match	160	3.45	93.3%	55.1%	Fails
Conditional redesign	320	2.44	99.8%	84.0%	Conditional

[PD] Usual variable burden = 100; incremental cost is 4.27 index points. Require 80% positive-direction detection at break-even under a two-sided 5% test, plus an 80% joint-power floor for economic and timely-recovery superiority. Economic clearance requires a one-sided 97.5% lower bound above fixed cost.

[PD] The response family mixes endpoint per-request policies; it is not observed compliance. Stress lowers presence to 0.20 / 0.60. Imperfect matching reduces pair-grade correlation from 1.00 to 0.50; matching shared risk never means sharing staff or capacity. Costs and remedy clocks stay fixed.

CONDITIONAL PLANNING IS NOT RECRUITMENT EVIDENCE

[PD] The redesign needs 640 independent domains and 25,600 original flags under the assumed matching and covariance. Verify availability, reference accuracy and variance in a blinded pilot. Separately power customer-wait, safety and ordinary-flow guardrails; these two gates do not certify them.

[PD] Joint floor = max(0, economic power + recovery power – 1). No endpoint independence is assumed. Variance bounds are exact for the finite law; normal confidence coverage remains approximate.

Price verified recovery rather than a withheld refund

[PD] Service burden is an assumed resource valuation, not a customer-payment balance. Under lower presence, avoided service burden is 0.7128 per flag; added variable work is 0.07128. Fixed cost 0.40 allocates setup 0.15, coordination 0.15 and production verification 0.10 once.

Avoidability a	h = 0.80	h = 0.90	h = 1.00
0.6	-1.38	-0.92	-0.46
0.8	-0.16	0.45	1.06
1.0	1.06	1.82	2.58

[PD] Cells are net index points; h is the non-overlap share. $Net_raw = a \times h \times 0.7128 - 0.07128 - 0.40$; $Net_index = 100 \times Net_raw / 9.3774$. Discount favorable service benefit only; new work and fixed cost stay fully charged. Changed prices require re-powering.

[PD] At full realization, fixed cost 0.20 / 0.40 / 0.60 yields net 4.71 / 2.58 / 0.44 index points under lower presence. These are cost sensitivities, not confidence intervals. Research-only adjudication is separate from recurring verification.

[PD] Later accepted-recovery probability changes the maturity estimand; revalue and re-power. With 1% unresolved receipt truth per arm, the service-contrast bound is 0.24 raw units, almost the 0.24152 stress margin. Retain original denominators and observed work.

KEEP THE THREE CLAIMS SEPARATE

[PD] A completed search, an accepted physical delivery and an avoided payment are different events. No refund cash is included. Future field value must price the signed mature outcome vector over original eligible volume, then subtract full incremental costs. No historical effect or cost input supports that estimate here.

[PD] Uniform service valuation is hypothetical. Use category-specific consequences if an additional late delivery creates customer burden rather than reducing it; customer acceptance and entitlement remain mandatory.

Buy evidence for the next bounded decision

[PD] Information value is independent of an assumed positive recovery effect. The states distinguish actionable recovery, a measurement-only problem and unresolved safe authority. The last state takes priority so the decision partition is exhaustive.

Next investment action	Recovery viable	Measurement only	Unresolved
Bounded recovery-only trial	15	-10	-12
Reconcile evidence / clock	4	9	-3
Hold investment; obtain authority	1	1	5

[PD] Priors are 0.40 / 0.35 / 0.25. Best prior utility is 4.00; perfect information can add at most 6.40. Signal accuracy q assigns each wrong signal probability $(1 - q) / 2$. All 27 signal-to-action policies are evaluated.

Signal quality q	EVSI	Evidence cost	Net EVSI
0.50	0.0000	0.50	-0.5000
0.70	2.5450	1.50	1.0450
0.90	5.1150	3.00	2.1150

A POSITIVE INFORMATION VALUE DOES NOT AUTHORIZE A HOLD

[PD] Utilities and likelihoods are hypothetical. Purchase only evidence that changes a permitted decision; its acquisition cost is distinct from later trial cost. Do not add this information value to service benefit, claim realized saving, or use it to bypass customer-remedy authority.

[PD] EVSI = expected posterior-optimal utility – prior-optimal utility; net EVSI subtracts acquisition cost. The proposed recovery-only trial preserves existing customer remedies.

Reproduce the decision boundary before valuing the mechanism

[AI] This case formalizes the Recovery–Remedy Evidence Contract (RREC): validate state at the decision instant, test incremental customer-accepted physical recovery, separate authority for remedy changes, and value mature outcomes over the original cohort.

Claim class	Meaning and replication boundary
[HR] Historical record	Statements and actions transcribed in the supplied project register. Original query results, workflow records and authored documents remain unrecovered. The register’s descriptions of direct evidence do not substitute for inspecting the originals.
[RC] Recalculated	Matched-population coverage, source arithmetic and scope reconciliation, tentative process screens, and exact constructed counterexamples. Count checks establish arithmetic consistency, not physical presence, implementation or causal impact.
[AI] Analytical inference	Last scan, current custody, feasible recovery and incremental value are different claims. A bundle effect does not identify the hold’s contribution; blank coding need not imply list exclusion; cadence alone does not identify event order.
[PD] Proposed design	RREC, state-validation sampling, a matched-domain recovery-only study, approximate power planning with finite-process checks, full-cost sensitivities and a separate information-value model. None represents a historical pilot.

Replication sequence

[PD] Freeze the workbook and exact selected cells. Reconcile matched and unmatched populations, blank mechanisms, period and geographic scopes, financial rates and the unresolved precedent versions. Preserve projections and targets as designs rather than observations.

[PD] Recover event and transaction histories, validate physical state and customer willingness, and define the original eligible cohort before assigning treatment. Independently reproduce the proposed law and calibrate presence, recovery chances, pairing quality, full costs and attainable recruitment.

THE DECISION THE CASE CHANGES

[AI/PD] A proposed refund hold becomes a component question rather than an assumed source of savings. Better detection and earlier physical recovery can be tested while the existing remedy timetable remains intact.

Public historical coverage is indexed to one matched category total = 100. Raw measurements, source identities, local paths, financial keys and adjacent-project evidence remain private.

RESEARCH REFERENCES / STATE, TIME AND CUSTOMER RECOVERY

Observe the state that matters and test the component that changes it

Primary research and official methods inform this case's state validation, component contrasts and prospective evaluation. They do not authenticate the unrecovered project originals or establish an operating result.

Jeffery, S. R., Garofalakis, M., & Franklin, M. J. (2006). Adaptive Cleaning for RFID Data Streams. Proceedings of VLDB, 163–174. Full author text. Physical movement and imperfect sensing can produce competing interpretations of observed reads. This is a measurement analogy, not a claim that the project used RFID; last-scan validity requires independent field evidence.

Snodgrass, R., & Ahn, I. (1985). A Taxonomy of Time in Databases. Proceedings of SIGMOD, 236–246. DOI 10.1145/318898.318921. Full author text. Valid time and transaction time are distinct. The case preserves physical event time, ingestion time and decision-time availability so later corrections cannot silently rewrite what the operating rule knew.

Hernán, M. A., et al. (2016). Specifying a target trial prevents immortal time bias and other self-inflicted injuries in observational analyses. Journal of Clinical Epidemiology, 79, 70–75. DOI 10.1016/j.jclinepi.2016.04.014. Primary abstract. Eligibility, treatment assignment and follow-up must align. Selecting only items that survived long enough for a recovery attempt would change the intended population.

Collins, L. M., Dziak, J. J., Kugler, K. C., & Trail, J. B. (2014). Factorial Experiments: Efficient Tools for Evaluation of Intervention Components. American Journal of Preventive Medicine, 47(4), 498–504. DOI 10.1016/j.amepre.2014.06.021. Primary indexed full text. Component contrasts and interactions answer questions a bundle comparison cannot. Any remedy-change component still requires separate authorization.

Tax, S. S., Brown, S. W., & Chandrashekar, M. (1998). Customer Evaluations of Service Complaint Experiences: Implications for Relationship Marketing. Journal of Marketing, 62(2), 60–76. DOI 10.1177/002224299806200205. Publisher abstract. Customers evaluate more than a final transaction; process and interaction matter. Physical receipt alone is not a complete customer-welfare endpoint.

Anderson, S. W., Baggett, L. S., & Widener, S. K. (2009 issue; 2008 online). The Impact of Service Operations Failures on Customer Satisfaction: Evidence on How Failures and Their Source Affect What Matters to Customers. Manufacturing & Service Operations Management, 11(1), 52–69. DOI 10.1287/msom.1070.0193. Publisher abstract. Perceived responsibility changes satisfaction relationships; no airline effect is imported into parcel recovery.

Access scope is recorded in the private research memo. No external study's effect size, financial result or certification is transferred to this project.

RESEARCH REFERENCES / UNCERTAINTY AND EVALUATION

Keep selection, shared resources and incomplete records inside the analysis

Primary research and official methods inform this case's state validation, component contrasts and prospective evaluation. They do not authenticate the unrecovered project originals or establish an operating result.

NIST/SEMATECH. e-Handbook of Statistical Methods, section 6.3.2.2: Individuals Control Charts. Official method. Comparable ordered observations support tentative natural-process limits. Unrun plans and targets are not observations; a crossing can prompt investigation but cannot identify the treatment's cause.

Laney, D. B. (2002). Improved Control Charts for Attributes. *Quality Engineering*, 14(4), 531–537. Primary publisher abstract. Dispersion-adjusted rate monitoring requires comparable numerator-denominator pairs. A source-mentioned chart without its subgroup data is not a reproduced control result, and a financial-policy change can alter the monitored construct.

Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20, 37–46. Primary publisher record. Independent adjudication and confusion tables support taxonomy reliability. Agreement on a last-scan label does not establish that the item was physically present or recoverable.

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. Worst-mover selection can generate apparent recovery without a treatment effect. The source lacks unit-level histories for an empirical correction; retain concurrent comparisons and preregister cohort selection.

Hudgens, M. G., & Halloran, M. E. (2008). Toward Causal Inference With Interference. *Journal of the American Statistical Association*, 103, 832–842. Primary manuscript. Extra recovery work can displace shared staff or transport. Define assignment domains around those dependencies and measure spillover rather than assuming parcel-level independence.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, 225, 200–230. Author manuscript. A later nonrandom staggered rollout needs group-time effects, valid comparison groups and explicit timing assumptions. A proposed randomized model does not validate an observational implementation automatically.

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. Evidence is valuable when it changes a decision. The separate illustrative decision model charges acquisition cost and does not add information value to the same operating benefit twice.

Access scope is recorded in the private research memo. No external study's effect size, financial result or certification is transferred to this project.

RESEARCH REFERENCES / MEASUREMENT AND OPERATING LEARNING

Make the customer outcome harder to obscure than the financial score

Primary research and official methods inform this case's state validation, component contrasts and prospective evaluation. They do not authenticate the unrecovered project originals or establish an operating result.

Collier, D. (2011). Understanding Process Tracing. *PS: Political Science & Politics*, 44, 823–830. Primary publication. Discriminating evidence supports contribution claims. Commissioning, analysis, authorship, approval and implemented operating ownership are separate propositions; workflow provenance limits the case's attribution grade.

Tucker, A. L., & Edmondson, A. C. (2003). Why Hospitals Don't Learn from Failures. *California Management Review*, 45(2), 55–72. Author manuscript. Local workarounds can leave systemic problems unresolved. More frequent exception recovery needs detection repair, accountable decision rights and feedback, not merely a larger queue of repeated touches.

The W. Edwards Deming Institute. The Funnel Experiment. Official exposition. Reactive adjustment to routine variation can increase instability. This case rejects causal claims from isolated weekly changes and distinguishes a proposed control plan from a demonstrated operating phase.

Campbell, D. T. (1976). Assessing the Impact of Planned Social Change. Original institutional record. Consequential indicators can alter behavior and recording. A refund gate directly changes the financial event being counted, so independent physical and customer outcomes must remain visible. This is a design critique, not an allegation of historical manipulation.

Goodhart, C. A. E. Problems of Monetary Management: The U.K. Experience. 1975 conference; revised volume listed as 1976 in the RBA archive. Targeting can change the relation between an indicator and the intended outcome. A reduced concession rate cannot alone validate a remedy-policy intervention.

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. Measures must connect to strategic decisions. An exposure calculation supports prioritizing validation; it does not replace an incremental-value estimate or confer decision rights.

OECD / European Commission JRC (2008). Handbook on Constructing Composite Indicators. Official handbook. Weights and aggregation affect interpretation. Expose customer outcomes and resource categories beside the proposed burden index, and show FMEA dimensions separately rather than treating their product as a probability.

Access scope is recorded in the private research memo. No external study's effect size, financial result or certification is transferred to this project.

Accept a bounded design contribution; keep the operating claim open

[AI] This is author-led, agent-assisted review, not external peer review, client approval or a consulting-firm ranking. Its verdict concerns source interpretation, reproducible analysis and a falsifiable design. It certifies no historical recovery effect or saving.

Major issue	Disposition in the revised case
State and temporal inference	Separated recorded scan, authenticated presence, feasible recovery and added value. Neither missing terminal records nor relative cadence establishes the refund-time state.
Population and visibility	Reconciled missing attribution and blank mechanisms against the matched parent population. Did not infer that missing codes make every alternative recovery signal blind.
Bundle and remedy policy	Separated enhanced recovery from refund delay. The exact two-world proof shows why a bundle result cannot establish the hold’s contribution; the proposed trial preserves the remedy clock.
Transport and economics	Kept the borrowed stage/site precedent, scope versions and unvalidated cost assumptions explicit. Charged additional work and gave no physical-service credit to withheld refunds.
Contribution and execution	Limited personal contribution to provisionally supported framing, directed analysis and design development. Preserved workflow authorship uncertainty, unassigned ownership and absent implementation.

Partner critique: can a COO act on page one?

[AI] Yes: assign a bounded state-and-timing validation and establish the recovery-only test’s owners, costs and customer safeguards. A separate policy owner must decide whether any refund change is permissible; the analytical author cannot supply that authority.

Minor revisions and remaining limits

[AI] Recover originals, reconcile precedent versions and calibrate the proposed outcome law, pairing quality, observation horizon, resource weights and feasible recruitment. Approximate planning results remain conditional; separate safety and customer-remedy checks require their own thresholds.

REVIEW DISPOSITION / PASS WITHIN STATED SCOPE

[AI] The case supports a proposed Recovery–Remedy Evidence Contract and bounded analytical leadership. It establishes neither historical savings nor a reason to delay a customer’s remedy.

Scope: internal evidence, methods, numerical and executive review. Acceptance is conditional on stated evidence boundaries; no external endorsement or firm ranking is implied.