

OPERATING CASE STUDY / CUSTOMER EXPERIENCE & LOGISTICS

Loss Prevention Portfolio Design

Turn a peak-loss investigation into controls with distinct benefits and accountable owners

An anonymized operating case on a peak-season loss investigation, prioritized control proposals and a subsequent management review. The analysis separates physical recovery from reporting changes and tests whether proposed benefits can be added.

PROJECT CONTEXT / LOSS DIAGNOSIS & OPERATING-ACTION DESIGN

BLUF / BOTTOM LINE UP FRONT

Fund distinct, preventable losses.
Verify recovery before claiming impact.

ANALYTICAL FINDING

Two actions share a source category. Assuming all other action sets are disjoint, their possible overlap changes the additive opportunity index of 100 to a conditional 77.1–100. Other overlaps remain unmeasured.

CONTRIBUTION IN THE SOURCE RECORD

Authored the diagnostic and action design; advanced a review mechanism. Provisional registry hoop support; operating execution remained with assigned owners.

EVIDENCE STATUS

Source-reconciliation and operating-design case; originals unrecovered. Proposed recoverability is not realized benefit; mature post-intervention loss is unmeasured.

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

EXECUTIVE ABSTRACT / DECISION AND OWNERSHIP

Authorize a loss-prevention pilot after the opportunity is reconciled

[AI/PD] The proposed Loss Prevention Contract (LPC) links a defined customer-loss population, distinct preventable events, an authorized control, and independent outcome verification. It makes the approval question explicit: which losses will this package prevent, at what service cost?

Abstract element	Decision-relevant synthesis
Context [HR]	A peak-season investigation became a prioritized operational proposal and management-review mechanism. The record contains changing headline series, overlapping recovery estimates and unverified implementation.
Objective [PD]	Prevent physical nonreceipt while keeping delayed recovery, displaced work and ordinary-flow harms visible. Avoid paying twice for the same recoverable event.
Approach [RC/AI]	Reconcile source versions, category arithmetic and opportunity intersections; distinguish document ownership from operating execution; design an independent field evaluation.
Findings [RC]	The category-rate arithmetic does not reproduce every reported share. Under a limited overlap scenario, the unique priority-action opportunity spans 77.1–100 against additive 100. Neither result measures recovered customer deliveries.
Decision value [PD]	Authenticate the source versions and action-level event sets, then fund bounded calibration. Test customer burden and intervention cost on the same original eligible population and maturity horizon.
Implications [PD]	The operating sponsor owns rollout; data stewards own definitions; site and technology owners activate controls; an independent evaluation lead validates mature outcomes and service guardrails.

EXECUTIVE ACTION AND THRESHOLD

[PD] Approve calibration, withhold an additive benefit commitment. Scale only when a lower confidence bound on independently verified benefit exceeds incremental cost and separately powered service and truth-quality conditions pass. Maintain a separate frame for ordinary parcels affected by new gates.

Basis: S1 selected register. [HR] historical record; [RC] recalculated; [AI] analytical inference; [PD] proposed design. No realized loss-prevention benefit is established.


EXHIBIT 1 / RECONCILE THE QUESTION

Preserve the loss-series fork before allocating responsibility

[HR] The source family changes its headline loss series between drafts. [RC] Within-version arithmetic also fails to reproduce every category share. [AI] A majority in-station contribution remains plausible under matched-rate assumptions; a supposedly stable exact share does not.

VERSION IDENTITY COMES BEFORE A RESPONSIBILITY SPLIT

Analytical reconciliation / different series cannot share a borrowed denominator



Source: S1 M01–M06/M14, D07. Diagram describes source reconciliation, not a measured effect.

Reconciliation test	Decision consequence
Headline population [HR/AI]	The earlier and later rates, counts and growth claims refer to different stated views. A plausible relabeling explanation remains an inference until lineage and exposure are authenticated.
Category-rate closure [RC]	Recompute each share against its own version total. Quoted category shares and rate-implied shares cannot be mixed to manufacture a complete decomposition.
Unexplained remainder [AI]	A residual is an unallocated part of the stated rate system. It does not automatically become upstream loss, scan failure or a recoverable operating opportunity.
Quantity and exposure [AI]	Identical rounded unit quotations do not establish an identical cohort. Different event definitions can still use the same shipped exposure; both identities need evidence.

RETAIN THE ROBUST DECISION

[AI/PD] Investigate station custody and handoff controls. Do not allocate funding using a precise boundary share until the source identity, component coverage and denominator contract reconcile.

Basis: S1 M01–M06/M14 and D07; recalculations/private version ledger. All public statements retain the conditional matched-rate premise.

EXHIBIT 2 / CUSTOMER OUTCOME

Separate physical nonreceipt from its operating record

[HR] The diagnostic describes genuine custody failures, missing delivery events and disputed boundary attribution.
[AI] The recorded loss label therefore combines customer harm with measurement behavior; correcting a label does not establish recovery of a package.

PHYSICAL OUTCOMES AND SYSTEM LABELS ARE DIFFERENT ENDPOINTS

Proposed measurement / one label can change through more than one pathway



Source: S1 M12/M14/M17/M28; proposed LPC. No customer count is inferred from package events.

Event or observation	What it can establish
Independent delivery or recovery proof [PD]	A defined receipt or recovery event, with identity, timing and source quality. Missing evidence requires uncertainty bounds; it is not automatic proof of physical loss.
Loss classification [HR/AI]	A system disposition generated from the scan history and classification rule. Its sensitivity and false-positive rate may change when the proposed controls change.
Customer concession [HR/AI]	A separately defined financial/customer-service event. Missing concession records do not prove absent customer harm, and classified loss does not prove a unique refund or customer.
Outcome maturity [PD]	Freeze the original episode and deadline, then wait the validated ascertainment window. Late source updates and later reclassification remain traceable.

[PD] Audit outcome truth independently of the dashboard whose performance is being evaluated. Maintain distinct counts for verified receipt, delayed recovery, physical nonreceipt and unresolved reference evidence.

WHY THE DISTINCTION CHANGES APPROVAL

[AI] A cleaner record can be valuable without preventing physical loss. Value that correction through better decisions and reduced administrative work; reserve customer-loss benefits for verified incremental outcomes.

Basis: S1 M01/M12/M14/M17/M28; Prentice1989 and proposed LPC measurement design. No source event count is equated to unique customers.


EXHIBIT 3 / DISTINCT OPPORTUNITY

Reconcile shared recovery events before adding action benefits

[HR] The priority list assigns recoverable loss to both early-flow containment and protected problem-solving capacity in the same source bucket. [RC] Their possible shared events change the unique opportunity even when both individual estimates are accepted for planning.

POSSIBLE OVERLAP REDUCES DISTINCT PLANNED OPPORTUNITY

Conditional scenario / additive priority-action opportunity = 100



Only the identified pair overlaps; every other event set is assumed disjoint.
Other overlap or interactions invalidate this limited floor. Not a confidence interval.

Source: S1 M18 / action narrative; conditional union calculation. No realized recovery inferred.

Scenario / additive proposal = 100	Interpretation
No overlap: 100 [RC/PD]	All claimed recoverable events are distinct and comparable. This is the favorable additive scenario, not a demonstrated property of the source.
Only identified pair fully overlaps: 77.1 [RC/PD]	Every event prevented by the smaller estimate is also prevented by the larger. The remaining actions are assumed disjoint from this pair and each other.
Other overlaps or interactions [AI]	The limited lower endpoint is no longer a portfolio guarantee. Common buckets do not prove identical events, and combined controls may change the attainable effect.

[PD] Create an event-level opportunity register with unique entity and episode keys, treatment versions, reachable mechanisms and mutually exclusive accounting rules. Compare single controls and the package where feasible; predefine interaction estimands.

THE INVESTMENT CONSEQUENCE

[RC/AI] Up to 22.9% of the additive opportunity is exposed to duplication in this one conditional pair scenario. Preserve the full uncertainty; do not replace an unsupported point estimate with an unsupported guaranteed floor.

Basis: S1 M18 and narrative Intervention. Conditional union arithmetic; no realized recovery, independently estimated effect or complete portfolio bound.


EXHIBIT 4 / PORTFOLIO ALLOCATION

Use projected entitlement to find coverage questions, not a spending verdict

[HR] The original analysis contrasts where loss was classified with where an existing program portfolio assigned projected value. [AI] These are different quantities. The juxtaposition motivates a coverage review; it cannot establish a funding shortfall or rank marginal returns.

DEFECT SHARE AND PROGRAM ENTITLEMENT ANSWER DIFFERENT QUESTIONS

Analytical distinction / neither is an observed allocation of incremental spending



Source: S1 M23–M24. Entitlement is not capital committed, expenditure or independently verified value.

Decision input	Evidence needed before reallocation
Outcome exposure [PD]	A reconciled, mature loss population by mechanism. A boundary label alone does not identify controllability or the value of preventing one event.
Projected program entitlement [HR/AI]	A model of potential value, with scope and validation status. It is neither capital committed nor actual program expenditure.
Reach and implementation [HR/PD]	Which event sets an active program can reach, when it becomes effective, and whether sites adopted it. Below-line and status-unknown items remain distinct from committed work.
Incremental cost and interaction [PD]	The cost of the next implementable action, its distinct preventable events, overlap with existing controls and any harm displaced to another workflow.

A BETTER PORTFOLIO QUESTION

[PD] Ask whether a material verified failure mechanism lacks an effective owner and tested control. Obtain costs and marginal effects before deciding how much to fund; a mismatch of shares is not a cost-effectiveness calculation.

Basis: S1 M23–M24, source validation questions and alternative portfolio totals. No entitlement amount is represented as expenditure or personal benefit.

EXHIBIT 5 / OPERATING CONTROLS

Pair each forcing function with evidence of recovery and a safe exception path

[HR] The revised proposal moves from staffing expectations to shift, dwell, return and dispatch controls. [PD] A control must preserve a route for unresolved work and authorized exceptions; blocking a transaction cannot make a physical package disappear from accountability.

CONTROL THE PHYSICAL EXCEPTION AND THE COST OF CONTAINMENT

Proposed LPC / a blocked transaction still leaves an accountable physical episode



Retain unresolved work, shifted queues and ordinary-flow delays.
A successful compliance check does not certify customer recovery.

Source: S1 D03; LPC proposed extension. Controls and safe exception paths require field validation.

Proposed control family [HR]	Required operating proof [PD]
Disposition before shift close	Physical location or unresolved custody ownership; carryover remains in the eligible frame rather than being closed administratively.
Dwell escalation and protected recovery capacity	Time-stamped intervention, usable capacity and recovery evidence. Retain work displaced from another queue and delays created by the control.
One verified return pathway	Identity-preserving receipt and storage proof, with a documented alternative during outages or a justified exception. Rejected transactions remain visible.
Staging-to-departure reconciliation	Physical package/manifest alignment, custody acceptance and independently reviewable exceptions. Track ordinary-flow delay as well as loss prevention.

TEST THE FAILURE MODE OF THE CONTROL

[PD] Pilot peak load, outages, unavailable approvers and invalid scans. Define stop, override and rollback authority before activation. An improved compliance number with more waiting or unrecorded work is not sufficient evidence of better service.

Basis: S1 D03 and implementation design. Grout2007; Strom etal2010; Little1961. External mechanisms inform proposed tests; no effect is transferred to this network.

EXHIBIT 6 / PLAN INTEGRITY

A reconciled plan is a constraint on execution, not evidence of results

[HR] The source family includes monthly targets, annual goals and action-specific recovery estimates. [RC/AI] These require different checks: rounding can explain a small component discrepancy, while a monthly target sequence is not an observed operating process.

PLANNED TARGETS REQUIRE A DIFFERENT TEST FROM OBSERVED PERFORMANCE

Source-contract review / do not turn targets into an operating time series

ROUNDING
Check compatible intervals

AGGREGATION
Use matching exposure weights

REALIZATION
Measure mature actual outcomes

A rounded component discrepancy may be compatible with the total.
Neither exact reconciliation nor a target trajectory demonstrates achievement.

Source: S1 M20–M22 and goal-table definitions. Plan checks do not establish operating SPC signals.

Plan-level check	Conclusion supported by the record
Annual component reconciliation [RC]	The displayed category totals do not equal the displayed overall total exactly. Their rounding intervals can be compatible; preserve the discrepancy and request unrounded values.
Monthly-to-annual aggregation [AI]	An annual exposure-weighted rate need not equal the simple average of monthly rates. Without matching monthly opportunities, neither weighting nor annual realization is established.
Seasonal target shape [HR/AI]	A sharp planned month-to-month change may reflect a strategy or an entry error. Keep it visible, authenticate the source, and do not smooth it into an invented observation.
Goal versus recoverability [PD]	A target is a management commitment; recoverability is a causal opportunity assumption. Compare them only on the same event definition, volume, maturity and feasible treatment.

RETAIN THE PLAN’S ACCOUNTING BOUNDARY

[HR/PD] A separate loss classification is excluded from the stated concession total. Adding it would change the metric definition and may count the same customer event again. Verify the mapping instead of enlarging the total.

Basis: S1 M20–M22/M28 and goal-table definitions. Mechanical target screens in the private appendix are not operating SPC evidence.

EXHIBIT 7 / DELIVERY AND ACCOUNTABILITY

Treat the management mechanism as the delivered asset and test operating adoption separately

[HR] The record attributes the diagnostic, action specifications and management-review design to the case author. It assigns execution to operating and program owners. [AI] A mature governance document can be a completed deliverable while the underlying controls remain unevaluated.

A DELIVERED ANALYSIS DOES NOT ESTABLISH ACTIVATED CONTROLS

Historical contribution and proposed evidence / assign each stage its own accountable role



Historical support is strongest for the diagnostic and review mechanism.
Activation, adoption and mature customer effect require separate records.

Source: S1 attribution, D08, implementation register. Authorship does not transfer program ownership.

Stage and accountable role	Evidence required or available
Analysis and design / case author [HR]	Provisional source support for authorship, revision and submission. Personal contribution is narrower than ownership of every referenced program or target.
Approval and prioritization / sponsor [HR/PD]	Explicit decision and accepted scope. Review requirements and named owners are evidence of governance; they do not establish activated controls.
Activation / site and technology owners [PD]	Versioned configuration, training, exception handling, verified eligible coverage and an auditable treatment start. An action date alone is insufficient.
Outcome assurance / independent evaluation lead [PD]	Frozen episodes, mature physical outcomes, retained harms, a valid counterfactual and incremental cost. Separate outcome effect from improved reporting or target attainment.

NEXT DECISION

[PD] Fund authentication and calibration before scaling. Pre-register a clustered evaluation of the complete package, require independent outcome truth, and make net benefit, service performance and unknown-outcome limits joint release conditions.

Basis: S1 master attribution, D08 and implementation register. Contribution: provisional registry hoop support; execution and mature outcome ownership not established.


TECHNICAL ANALYSIS

A prevention control must change custody, not merely its record

[AI/PD] The proposed Loss Prevention Contract (LPC) separates physical prevention, administrative recording and customer burden. The graph makes an apparent reporting improvement distinguishable from a tested operating effect.

PREVENTION AND RECORD RECLASSIFICATION HAVE DIFFERENT CAUSAL PATHS

Proposed DAG / latent physical burden and its independent measurement remain distinct



Y is determined by L and D; O measures Y with reference quality J. The planning model assumes valid ascertainment.

[PD] Z assigns the package; A denotes activated controls and M the physical recovery mechanism. L is true unresolved physical loss; D is costly delayed or reworked recovery; together they determine true mature burden Y. Controls can change both L and D. Their customer value cannot be inferred from compliance alone.

[PD] Q is capture and classification; R is recorded Lost. A can change Q and hence R without changing L. L also influences R. Independent ascertainment J, together with L and D, produces measured outcome O. The planning model assumes valid ascertainment; audit disagreement or missing truth requires sensitivity rather than automatic certainty.

[PD/AI] C names season, volume, weather, case mix and capacity; U includes unrecorded changes and concurrent programs. These can affect activation, recovery, delay and recording. The registered analysis supports candidate mechanisms and control specifications, not historical identification of their effects. Residual differences may arise through these unmeasured paths. Assignment identifies an ITT effect only within the proposed trial population.

Basis: registered diagnosis and proposed controls [HR]; LPC causal and measurement graph [PD]. Arrows are hypotheses, not demonstrated project effects.

TECHNICAL ANALYSIS

Freeze the loss construct before estimating what can be prevented

[HR/AI] The source uses physical custody, operational loss labels and concession events in overlapping ways. Treating them as interchangeable would change both the problem size and the meaning of success.

Quantity	Contract and interpretation [AI/PD]
Recorded-loss rate	Unique qualifying recorded events divided by the stated eligible exposure, with population, period, cutoff and attribution version frozen. This is descriptive unless treatment is identified.
Physical-loss risk	True unresolved physical disposition by the prespecified horizon among an upstream fixed cohort. Independent investigation measures it imperfectly; an absent scan is not sufficient truth.
Customer payment or concession	A separate recognized event with its own timing and eligibility. No corresponding concession does not prove no customer harm, nor does a loss label prove a refund.
Action-package ITT	Mean mature burden under usual practice minus mean under randomized package assignment. Positive change is benefit; activation and actual execution remain mediators.

[AI] Distinct event sets can share the same exposure denominator. Conversely, similarly named rates can have different populations or maturity windows. The recorded fraction without a corresponding concession does not establish the share of true physical loss, universal refund incidence or a necessary denominator difference.

[PD] For a binary fixed-cohort example, recorded-loss probability equals $S_e p + (1 - S_p) \times (1 - p)$, where p is true physical-loss risk and S_e/S_p are ascertainment sensitivity/specificity. Changing either measurement property can change the dashboard while p stays fixed. This is a measurement identity under its assumptions, not an estimated local error rate.

[PD] Preserve original promise, custody eligibility and follow-up through reroute, hold, cancellation and return. Maintain event-time and recognition-time views separately. Reclassifying an upstream event may improve accountability without preventing the event; independently verify the physical endpoint before claiming customer benefit.

[AI] A historical before/after count from another team’s fix concerns a different intervention and unequal seasonal windows. It is neither this author’s delivered benefit nor a causal estimate for the newly proposed package.

Basis: source population and attribution conflicts [HR]; measurement identity and estimands [AI/PD]; Prentice; VanderWeele. No historical sensitivity or causal effect is estimated.

TECHNICAL ANALYSIS

Preserve a robust priority without manufacturing a precise decomposition

[RC/AI] In-station loss remains the larger reported boundary in the examined cuts, but the repeated headline share does not reconcile exactly to the component rates. The priority is more defensible than its apparent precision.

Reconciliation question	Required treatment [AI/PD]
Which series and version?	Keep operational and concession-related cuts separate. Select the version appropriate to the stated question; do not average them or label a later filename as verified truth.
Does the boundary bridge close?	Recalculate component/parent shares on the same basis. Retain the unassigned residual; it is not automatically upstream loss or a measured miscellaneous cause.
Are driver definitions stable?	Crosswalk pickup, return and reprocess events. An unchanged label can conceal a changed mechanism or eligibility rule. Direction disagreement cannot be explained by one common positive rescaling.
Do target rows reconcile?	Assess rounding intervals and hierarchy before asserting an arithmetic error. A nearly closing target table does not establish independent outcome categories or attainable benefits.

[HR/AI] The source presents both a boundary decomposition and a separate problem-solving classification. Their named components exceed the parent when naively combined. This can reflect overlap, mismatched populations, versions or definitions; it does not identify which records overlap. Require a record-level crosswalk before turning classifications into additive opportunity pools.

[PD] Construct a frozen unit-by-period ledger with physical identity, custody boundary, event code, attribution version, recognition time and source lineage. Assign an exclusive primary class only when the policy defines one, retaining secondary contributing mechanisms. An analytical partition is an accounting convention, not proof of causal origin.

[AI] Common positive scaling preserves rank and direction within the same observations. Where cuts disagree on direction, examine changing membership, definitions, attribution and windows. The case cannot choose the more dramatic trend merely because it supports the narrative.

[PD] Recompute the decision under each admissible version. If both prioritize the same boundary, state that limited robustness; if the proposed control depends on a disputed driver, hold that specific release pending reconciliation.

Basis: versioned source discrepancies and recalculation [HR/RC]; private ledger retains raw alternatives. Public conclusion does not certify the exact reported split.

TECHNICAL ANALYSIS

Specify the control’s failure and exception paths before blocking work

[HR/AI] The action revision shifted from reminders and staffing commitments toward forcing functions. Several proposed controls still depended on staff, system changes or authorized discretion; a design label did not remove those dependencies.

Control class [HR/PD]	Acceptance evidence and failure mode [PD]
Shift-close or disposition gate	Every eligible item has an independently valid disposition; test false closure, blocked legitimate completion, unavailable systems and queue transfer.
Dwell escalation and recovery	A timed trigger reaches a staffed authorized responder, who can locate or recover the item; escalating without capacity only moves the queue.
Protected staffing coverage	Verify scheduled and delivered recovery capacity, competing work and authorized exceptions. A staffing lock is not a physical interlock.
Return-path or departure validation	Reconcile actual custody with the transaction or manifest; test valid alternate paths, offline conditions, late arrivals and conflicting records.

[PD] A release contract specifies trigger, protected population, evidence required, control action, legitimate exception authority, response deadline and verification. A block on an invalid state should not become an unreviewable block on all work. Record exception reason and subsequent outcome; an undocumented bypass fails the contract.

[AI] Grout’s original design guidance treats the prevention device as a potential source of error. A clinical randomized hard-stop study also provides a counterexample to equating changed behavior with improved ultimate outcomes. These sources motivate testing control failure and delay here; their clinical effect sizes do not transfer.

[PD] Test controls on representative normal, abnormal and degraded-system scenarios before field assignment. Quantify false blocks, unsafe releases, lost traceability and recovery workload. Maintain authorized rollback and interruption handling. Restrictions must be approved by the system and operating owners; the case itself grants no authority.

[AI] Little’s flow relation applies under a stable matched boundary, not automatically during peak disruption. Closing a dashboard item cannot itself remove physical unfinished work. Check whether the gate reduces unresolved custody or merely shifts the waiting location.

Basis: registered action specifications [HR]; prospective acceptance design [PD]; Grout; Strom et al.; Bainbridge; Little. Control activation and field performance remain unestablished.

TECHNICAL ANALYSIS

Value the incremental package rather than adding every opportunity

[AI] A defect share, an estimated recoverable amount and a financial entitlement answer different questions. None is a budget allocation, an approved commitment or an independently identified package effect.

[RC/AI] A sensitivity that varies only the overlap of two action estimates within one parent bucket puts their combined contribution below or equal to their naive sum. Holding every other action disjoint gives a total between about 77 and 100 on the nominal action-sum index. This is a conditional overlap illustration, not a credible interval or a lower bound for the whole portfolio.

[PD] If recoverable sets E1 and E2 refer to identical units, horizon and outcome, their union is $|E1| + |E2| - |E1 \cap E2|$. Common bucket membership alone does not identify the intersection. The set calculation additionally assumes prevention is nonnegative and controls do not create or displace failures. Actual package effects require treatment contrasts.

Value layer	Decision rule [AI/PD]
Opportunity entitlement	Document the expected benefit construct, owner, scope, maturity and validation status. It is not spending or guaranteed recoverability.
Incremental package effect	Compare the package against the operating baseline including prior fixes and concurrently funded work. Deduplicate affected units and avoid benefit already realized elsewhere.
Net operating value	Price the identified outcome over the same eligible volume and horizon; subtract implementation, recovery capacity, maintenance and exception handling once.
Decision information	Value whether additional evidence changes the selected plan. Keep EVSI separate from the operating benefit of the same prevented loss.

[AI/PD] An interaction is scale-specific. For burden means μ_{ab} under two controls, $I = \mu_{11} - \mu_{10} - \mu_{01} + \mu_{00}$ measures departure from additivity on that burden scale. Positive I means the joint burden reduction is smaller than the sum of isolated reductions. Shared beneficiaries, displaced capacity or complementarity can produce interaction; component budgets cannot diagnose it.

[PD] A later factorial extension could estimate the interaction if components can be assigned independently. The current two-arm package trial does not identify it. Freeze the scale and multiplicity plan and power the interaction separately; an imprecise nonsignificant interaction does not validate addition.

Basis: source opportunity estimates and overlap uncertainty [HR/RC]; conditional union and potential-outcome algebra [AI/PD]; Kahan et al. No estimated recovery or actual funding share is claimed.

TECHNICAL ANALYSIS

Investigate extremes without mistaking seasonality or coding for a signal

[HR/AI] The work began from high-loss regions during peak disruption. Selected extremes, seasonal exposure and changing attribution constrain any inference from their subsequent movement.

Variation and selection

[PD] MR is the absolute adjacent change. Tentative individual limits are $\text{mean}(x) \pm 2.66 \times \text{mean}(\text{MR})$; the MR upper limit is $3.267 \times \text{mean}(\text{MR})$. Short, seasonal or definition-changing sequences are non-diagnostic. Plans and source versions are not achieved operating trajectories.

[HR/RC/AI] The private audit retains 39 screens: 14 sparse historical comparisons, 11 source-version comparisons, six plans, four baseline-target pairs, one censored-bound sensitivity and three artifact comparisons. All historical points lie within provisional limits; plan-shape crossings are not operating signals. No stable phase, mature package-outcome panel or compatible repeated numerator/exposure panel supports a Laney p-prime chart. Two-point screens are non-diagnostic by construction.

[PD] Freeze selected regions or repeat-event cohorts before follow-up. Under an explicitly hypothetical normal ranking model, selecting the worst tenth gives mean 1.755 SD. Assumed repeat reliability 0.3, 0.6 or 0.9 implies untreated regression of about 1.23, 0.70 or 0.18 SD. This parameterizes the selection risk; it does not estimate how much of this project's movement was reversion.

Independent measurement and moderation

[PD] Sample a random core of 240 custody histories, plus flagged difficult cases, from a declared upstream frame. Two blinded coders classify physical disposition, recorded status, boundary, delay and uncertainty. Keep independent labels, adjudicate disagreements and report confusion matrices, prevalence, agreement and kappa intervals. Weight oversampled classes; cluster by shared route or source family.

[PD] The iid worst-case raw-agreement normal-screen margin is about ± 6.3 percentage points at 95%; it is not kappa precision. Simulate kappa precision using pilot prevalence and dependence. Audit false blocks and valid exceptions as well as missed losses.

[PD] Prespecify business-line, node-type, custody model, volume stress and recovery-capacity moderators. A control requiring physical presence cannot solve a purely remote attribution error. Recovery without spare capacity may substitute delay for loss; incompatible scan paths require a separate treatment version.

Basis: selected source cohorts and measurement limits [HR]; prospective audit and regression sensitivity [PD]; NIST SPC; Laney; Barnett et al.; Cohen.

TECHNICAL ANALYSIS

Attribute the analysis and action design without importing successor results

[HR/AI] The register provides a substantial authorship and review trail for the deep dive, action sizing and governance design. Approval, technology delivery, station execution and mature benefit remain separate contribution claims.

Claim	Process-tracing grade and boundary [HR/AI]
Led the diagnostic document	Provisional hoop: registered tasking, submission and document metadata support a bounded author role; originals and independent contemporaneous authentication remain necessary.
Reframed and sized priority actions	Provisional hoop for the registered revision and sizing. Executive direction also shaped the forcing-function doctrine; estimating opportunity is not delivering recovery.
Created the recurring review structure	Provisional hoop for a registered creator role and owner/status tables. Co-editors and workstream owners retain their respective contributions.
Owned every target and program	Not supported: glidepath and external program sizing came from other contributors; some validation requests remained open.
Delivered the reported technical fix or later loss reduction	Not established for this project. The fix belonged to another team; successor projects and recurring files are not this package’s outcome evidence.

[AI] Repeated wording is not independent corroboration. One claimed launch contribution nearly reconstructs the entire increment implied by the reported current level and growth rate; this arithmetic does not isolate a launch effect. A filename or copied status is not a smoking gun for execution. No doubly decisive evidence identifies the design-to-customer-effect chain. Missing completion records do not prove no work occurred elsewhere.

Consolidated limitations and reproducibility

[HR/AI] Measurement is constrained by unresolved physical versus recorded loss, taxonomy, cohort maturity and nonclosing decompositions. Identification lacks concurrent controls and intervention fidelity. Transport is limited by selected regions, seasonal stress and custody models. Contribution relies on registry transcriptions without original-source authentication. Costs, overlaps, uptake and opportunity validation are insufficient for historical economics.

[PD] Reproduction requires frozen source versions, cohort and entity dictionaries, record-level overlap crosswalks, maturity extracts, control/exception logs, blinded outcome labels and executable calculations. The private annex retains raw identities and values. Public indexing preserves only interpretable comparisons; it does not cure a definition mismatch.

Basis: derivative authorship and governance evidence [HR]; process-tracing grades and limitations [AI]. Supported case type: diagnosis and operating design, not causal impact.

TECHNICAL ANALYSIS

Test the package on a fixed exception cohort and monitor displaced harm

[PD] Randomize independent custody/resource domains to the LPC package or usual practice. Freeze exception eligibility before any control can hold, divert or relabel an item; the planning model is wholly hypothetical.

[PD] Follow every original episode to 42 days after its original delivery promise. Define disjoint dispositions: independently adjudicated unresolved physical loss; physically recovered but delayed or reworked; ordinary verified receipt. Preserve routing and cancellation histories. An absent administrative scan is insufficient to establish physical truth; unresolved adjudication requires prespecified bounds.

[PD] The primary estimand is usual-minus-package ITT in mean cost-weighted mature burden per original eligible exception. Positive change is benefit. Physical-loss risk and costly recovery are separate mandatory outcomes. Assignment changes package availability and execution opportunities; comparing compliant with noncompliant episodes does not identify an execution effect.

[PD] The model proposes 80 exception episodes per independent domain after an initial smaller design. Shared domain risk induces correlation; episode count cannot substitute for independent domains. The package converts some otherwise unresolved losses into costly recovery. It assumes no additional ordinary-flow harm and valid independent ascertainment. Those assumptions require pilot testing, not retrospective confirmation from this case.

[PD] The model pages size directional benefit testing at alpha 0.025 and economic lower-bound clearance under an explicit finite probability law. The efficacy boundary changes the distribution; do not reuse only the full-effect variance. Compare MDE and break-even on the same fixed reference burden and service volume. Cost is allocated per eligible episode over the same horizon, not multiplied by research recruitment.

[PD] Independently monitor the ordinary-flow population that never entered the exception cohort: false holds, missed promises, congestion and unsafe releases can occur there. Preregister noninferiority margins and precision separately. Cluster shared queues, staff and systems together; where this leaves too few independent domains, redesign and re-power rather than counting correlated episodes as new clusters.

[PD] Parallel assignment preserves concurrent controls through outcome ascertainment. If staged rollout is required, estimate group-time ATT against not-yet-treated domains with no anticipation, common support and interference controls. Use baseline and event-time diagnostics, including equivalence checks; nonsignificant pretrends are not proof. Recalculate power for that design before relying on it.

Basis: prospective randomized field design and synthetic planning law [PD]; a trial is not executed. Staggered alternative follows Callaway–Sant’Anna and has no power claim here.

TECHNICAL ANALYSIS

Four propositions connect error prevention to outcomes that matter

[AI/PD] LPC is a proposed framework: establish a defensible loss construct, a reachable physical mechanism, an authorized control, incremental portfolio value and independently measured customer outcomes.

Proposition [PD]	Disconfirming evidence or boundary [PD]
P1: Reachable physical controls reduce mature loss in their eligible cohort.	A precise failure to reach the prespecified minimum physical benefit, or excess service harm, contradicts the claimed advantage.
P2: Recording-only changes cannot deliver physical prevention without another mechanism.	Reproducible physical benefit under an intervention demonstrably limited to recording would challenge the proposed exclusion restriction.
P3: Shared beneficiaries or constrained resources can make package value nonadditive.	A powered factorial study can bound the interaction near zero on the prespecified scale; common bucket labels alone cannot decide it.
P4: An authorized exception path preserves necessary flow while retaining control accountability.	Unsafe release, persistent bypass or excessive delay contradicts the specified tradeoff; all deviations and outcomes remain auditable.

[AI] Error-proofing research supports redesigning the process rather than relying solely on reminders, while requiring the device itself to be tested. The source’s control specifications are consistent with that design logic; this case does not verify their effectiveness. Automation can leave operators with difficult exceptions, so authority and recovery capacity remain part of the intervention.

[AI] Surrogate-endpoint research explains why a favorable status metric does not generally identify a favorable customer effect. LPC extends the practical measurement contract by retaining separate physical, recording and service-burden outcomes. It does not claim a formal surrogate paradox occurred historically.

[AI] Deming’s variation discipline opposes responding to unqualified noise as a proven process shift. Campbell and Goodhart motivate monitoring whether pressure on recorded loss changes coding or blocking. Ittner cautions against equating measurement architecture with performance; composite design makes the value placed on delay versus loss explicit.

[AI/PD] Tucker and Edmondson distinguish immediate recovery from learning that changes recurring conditions. A review forum is a mechanism to test, not a result. Preregister the meaningful effect and harm margins. An upper confidence bound below required benefit contradicts it; an interval spanning the threshold is inconclusive. Economic success cannot waive unresolved safety or truth requirements.

Basis: proposed LPC propositions [PD]; Grout; Strom et al.; Bainbridge; Prentice; VanderWeele; Deming; Campbell; Ittner et al.; OECD; Tucker and Edmondson. No historical effect confirmed.

PROSPECTIVE MODEL / LOSS PREVENTION CONTRACT

Count recovered parcels, not recoded records

[PD] Propose a randomized field pilot of the full Loss Prevention Contract (LPC) versus usual controls in independent custody domains. Every numerical input is hypothetical; no trial or field effect is observed.

[PD] Freeze custody-exception eligibility before assignment. Retain every original episode through rerouting, transfer, withdrawal and recovery. Independently adjudicate physical disposition at the original promised date plus 42 days; validate that horizon in the pilot. A missing administrative scan is insufficient evidence of physical loss.

Mutually exclusive endpoint	Resource-proxy loss	Meaning
Confirmed unresolved physical loss	20	Independent custody/disposition evidence
Delayed or reworked recovery	2	Physical receipt established; recovery cost retained
Ordinary verified receipt	0	No counted physical loss or recovery burden

[PD] Primary ITT is usual-minus-LPC mean mature customer burden per original exception episode. Report physical-loss risk separately. Unknown truth remains in the cohort with conservative outcome bounds; the planning law assumes complete adjudication.

PROTECT THE POPULATION OUTSIDE THE TARGET

[PD] Physical-loss risk must pass a separately powered gate. Monitor ordinary-flow parcels for new holds, delay, damage and displaced work. Exception-cohort power does not size these guardrails or establish an all-parcel network effect.

[PD] Isolate custody, return-flow and staffing dependencies across domains. Randomize the offer; do not condition the effect on execution. Loss weights and the 42-day horizon are planning assumptions.

PROSPECTIVE MODEL / LOSS PREVENTION CONTRACT

Recovery still consumes resources when prevention works

[PD] The planning mechanism moves some would-be physical losses into delayed recovery. It does not turn every recovered parcel into a costless on-time delivery.

Hypothetical domain law	Usual	Full LPC
Physical-loss risk: lower / higher-risk domain	0.15 / 0.21	0.09 / 0.126
Delayed-recovery risk: lower / higher	0.20 / 0.20	0.26 / 0.284
Would-be loss converted to recovery	0%	40%
Mean burden, usual = 100	100.0	67.6

[PD] The two domain types are equally likely. Outcomes are independent within a type; the shared risk creates dependence within each domain. The efficacy path converts a fraction $0.40 \times \lambda$ of would-be losses into recovery. Under common random draws, increasing λ changes only loss 20 to recovery 2, proving stochastic ordering.

[RC/PD] A separate label-only counterfactual improves specificity at fixed physical states. Recorded Lost falls 17.51%, but modeled physical-loss risk and customer burden do not change. Better records may have value; they are not recovered parcels or an additional randomized arm.

THE LAW HAS A BOUNDARY

[PD] The planning law assumes no new harm to otherwise ordinary flow. If the package creates extra delay, loss or displacement, re-estimate the joint outcome law and economics; a changed code cannot conceal that harm.

[PD/RC] All risks are conditional on synthetic custody-exception eligibility, not historical rates. Independent ascertainment and blinded reference adjudication are required; oracle uncertainty is not powered away.

PROSPECTIVE MODEL / LOSS PREVENTION CONTRACT

Detect break-even benefit before treating the plan as investable

[PD] Require 80% power for benefit over zero at one-sided alpha 0.025, and 80% probability that a 97.5% lower bound exceeds break-even at the full planning effect.

[RC/PD] Modeled burden reduction is 32.4 reference points. Net resource value = reduction × avoidable share a × non-overlap h – cost 12. Base a = h = 1 gives break-even 12; no cash is estimated.

Domains / arm, total	Episodes / domain	80% MDE	Economic power	Both gates
24 / 48	20	Outside family	36.4%	Fail
71 / 142	80	11.98	>99.9%	Conditional
170 / 340	80	7.78	>99.9%	Conditional

[RC/PD] More episodes cannot replace independent domains. Base minimum: 71 per arm. Higher heterogeneity requires 141; candidate 170 clears that stress. Recruitment remains unestablished; calibrate first.

Net at full effect / cost	a = 0.50; h = 0.80	a = 0.75; h = 0.90	a = h = 1.00
9	3.96	12.87	23.40
12	0.96	9.87	20.40
15	-2.04	6.87	17.40

[PD] Cost 12 allocates technology 5, staffing/support 4 and verification 3 per fixed eligible service volume. Research sample size does not change this allocation. Recovery effort is already in loss; research-only auditing is separate. Shares apply uniformly to the net contrast; otherwise reprice each outcome.

[PD] Future cash = eligible exception volume × (mature loss-risk ITT reduction × avoidable loss cost + recovery-risk ITT reduction × avoidable recovery cost) × non-overlap – incremental cost. Increased recovery makes its reduction negative. Field inputs are unavailable.

[PD/RC] Finite-law inference is conditional on the specified nuisance distribution. Independent-domain availability, truthful ascertainment, guardrail power and field calibration are not established by these calculations.

PROSPECTIVE MODEL / LOSS PREVENTION CONTRACT

Buy evidence that changes the portfolio decision

[PD] This separate information model values resolving scope and overlap before committing to a pilot portfolio. It does not value assumed customer improvements.

[PD] States are comparable disjoint prevention sets, materially overlapping sets on a common basis, and incompatible scope or clock, with priors 0.40 / 0.35 / 0.25. Incompatible basis takes precedence when problems coexist; its reconciliation plan includes necessary overlap work.

Planning action / state	Disjoint	Overlap	Incompatible
Prepare bounded full-package pilot	16	-15	-30
Remove duplicated opportunity first	5	14	-12
Reconcile scope before sizing	-3	-4	12

[RC/PD] Prior-optimal overlap reconciliation has utility 3.90; perfect information adds at most 10.40. A three-label signal is correct with probability 0.50, 0.75 or 0.90; wrong labels share the remainder equally. Choose the highest posterior expected utility.

Signal quality	EVSI	Evidence cost	Net EVSI
Low	0.62500	0.50	0.12500
Base	4.78125	1.50	3.28125
High	8.15250	3.00	5.15250

[PD] Utilities already include subsequent planning costs; subtract evidence acquisition once. Do not add this information value to the trial benefit or treat a modeled opportunity as approved funding. Source/authority, customer validity and safety gates remain mandatory.

THE INVESTMENT DISCIPLINE

[AI/PD] Prioritize the distinct preventable customer problem, then verify the physical result. A cleaner loss code, a recoverability estimate and an accountable owner each support a different decision.

[PD/RC] Hypothetical decision utilities, priors and signals. No historical action effect, approved investment, realized savings or final customer improvement is inferred.

Reproduce the decision before treating the opportunity as an outcome

[AI] This is a source-reconciliation and operating-design case. Its strongest contribution is an actionable diagnostic and control proposal, with a measurable path to validate customer value.

Claim class	Meaning and replication boundary
[HR] Historical record	A statement or attributed action in the supplied project register. Original operational artifacts remain unrecovered; transcription accuracy is separate from source authentication.
[RC] Recalculated	Compatible arithmetic, conditional set bounds, rounding intervals and explicitly hypothetical planning results. Recalculation does not make a source estimate empirically valid.
[AI] Analytical inference	A conclusion with stated premises: a population fork, uncertain attribution, incomplete benefit additivity or a recording path that bypasses physical recovery.
[PD] Proposed design	The LPC framework, prospective trial, evaluation thresholds, economic scenarios and information-value choices. These are not implemented controls or observed effects.

Reproduction sequence

[PD] Verify the frozen source hash and selected cells. Preserve each series version and its units, periods, categories and exposure. Recompute component shares and rounding compatibility. Build the recoverable-event crosswalk before substituting a union for the source’s additive estimate.

[PD] Reconstruct the independent domain law, monotone prevention mechanism and retained recovery burden. Recalculate directional power, the economic boundary, all smaller allocations and the separate information model. Test the measurement-only counterfactual without treating changed classification as physical prevention.

WHAT THE ANALYSIS CHANGES

[AI] The short list is an opportunity hypothesis. It becomes an investable prevention package when distinct reachable events, authorized controls, independent physical outcomes, service harms and incremental costs share one auditable definition.

Labels: [HR] historical record; [RC] recalculated; [AI] analytical inference; [PD] proposed design. Source basis S1; see private annex.

RESEARCH REFERENCES / METHODS AND INTERPRETATION

Design controls around physical outcomes and credible exception handling

Primary research and official methods support the interpretation and proposed tests. They do not authenticate the historical project or supply its missing outcome data.

Grout, J. (2007). Mistake-Proofing the Design of Health Care Processes. AHRQ Publication 07-0020. Full original report. Control reliability, displaced errors, false alarms and obstructed flow inform the proposed failure-mode tests. The case imports no health-care effect estimate.

Strom, B. L., et al. (2010). Unintended Effects of a Computerized Physician Order Entry Nearly Hard-Stop Alert to Prevent a Drug Interaction: A Randomized Controlled Trial. *Archives of Internal Medicine*, 170(17), 1578–1583. Primary indexed abstract. Behavioral compliance and downstream delays illustrate why a stronger block requires outcome and service evaluation.

Prentice, R. L. (1989). Surrogate Endpoints in Clinical Trials: Definition and Operational Criteria. *Statistics in Medicine*, 8(4), 431–440. Publisher abstract. Additional conditions are needed before a proxy comparison identifies the intended endpoint. A loss-record change is not assumed to be a validated customer-outcome surrogate.

VanderWeele, T. J. (2013). Surrogate Measures and Consistent Surrogates. *Biometrics*, 69(3). DOI 10.1111/biom.12071. Publisher abstract and indexed author-manuscript sections. A favorable proxy effect and outcome association do not generally guarantee a favorable outcome effect. The case makes no historical claim of a formal surrogate paradox.

Kahan, B. C., et al. (2025). Guidance for Protocol Content and Reporting of Factorial Randomised Trials: Explanation and Elaboration of the CONSORT 2010 and SPIRIT 2013 Extensions. *BMJ*, 388, e080785. Full author-repository article. Package, component and interaction effects require distinct estimands; the proposed two-arm power does not establish factorial interaction power.

Bainbridge, L. (1983). Ironies of Automation. *Automatica*, 19(6), 775–779. Publisher abstract. Difficult exceptions can remain with operators when routine actions are automated. The proposed design therefore specifies authority, evidence and recovery capacity for exceptions; it infers no historical workforce behavior.

Little, J. D. C. (1961). A Proof for the Queuing Formula: $L = \lambda W$. *Operations Research*, 9(3), 383–387. Publisher abstract. Under the relevant stable-system conditions, queue stock, flow and residence time are linked. Closing a record does not physically remove unfinished work; historical peak stationarity is not established.

Private research memo records access scope: full text, primary abstract or continuing reference verified earlier in the series. No external effect size is transferred.

RESEARCH REFERENCES / METHODS AND INTERPRETATION

Separate compatible arithmetic from identifiable operating effects

Primary research and official methods support the interpretation and proposed tests. They do not authenticate the historical project or supply its missing outcome data.

NIST/SEMATECH. e-Handbook of Statistical Methods, §6.3.2.2: Individuals Control Charts. Official method. The case computes tentative limits where pairs or short sequences exist while separately testing comparability. Targets, source revisions and censored bounds do not become a stable operating process.

Laney, D. B. (2002). Improved Control Charts for Attributes. *Quality Engineering*, 14(4), 531–537. Publisher abstract. Extra-binomial variation requires actual compatible counts and exposures. Large source volumes and rounded loss rates do not provide a valid repeated-rate panel.

Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20, 37–46. Publisher record. Independent classification uses agreement, confusion matrices and kappa with uncertainty. Shared coding error can produce agreement without correct physical-outcome truth.

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. Selecting the worst regions can generate apparent recovery without treatment. This case parameterizes the risk rather than estimating it from absent untreated repeats.

Collier, D. (2011). Understanding Process Tracing. *PS: Political Science & Politics*, 44, 823–830. Primary publication. Authorship, approval, implementation and outcome contribution require different evidence tests. A reported metadata trail in a derivative register remains weaker than authenticated underlying artifacts.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, 225, 200–230. Author manuscript. Group-time effects inform the alternative staggered rollout when randomization is unavailable; they do not inherit the primary randomized design's power or assumptions.

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. The case derives its own finite-state information model, without using this article's regression estimator or adding information value to operating benefit.

Private research memo records access scope: full text, primary abstract or continuing reference verified earlier in the series. No external effect size is transferred.

RESEARCH REFERENCES / METHODS AND INTERPRETATION

Link measurement, learning and governance without overstating benefit

Primary research and official methods support the interpretation and proposed tests. They do not authenticate the historical project or supply its missing outcome data.

Tucker, A. L., & Edmondson, A. C. (2003). Why Hospitals Don't Learn from Failures. *California Management Review*, 45(2), 55–72. Author manuscript. Local recovery and system learning are different mechanisms. A recurring review can preserve ownership, but the forum itself does not establish sustained loss prevention.

The W. Edwards Deming Institute. The Funnel Experiment. Official exposition. Adjusting to routine variation can worsen performance. In this record, source definitions, target trajectories and maturity must first be separated from actual comparable observations; no historical tampering effect is identified.

Campbell, D. T. (1976). Assessing the Impact of Planned Social Change. Original institutional record. Consequential indicators can distort behavior. The proposed LPC retains original episodes and independent physical outcomes so classification, closure or selection cannot silently substitute for customer recovery.

Goodhart, C. A. E. Problems of Monetary Management: The U.K. Experience. 1975 conference; revised volume listed as 1976 in the RBA archive. The measurement critique informs the separation of targets, operational labels and physical outcomes. The case does not allege historical manipulation or infer intent.

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. More detailed measures and a formal review cadence do not, by themselves, establish a causal performance effect.

OECD / European Commission JRC (2008). Handbook on Constructing Composite Indicators. Official handbook. Weighting and compensation require explicit choices. The LPC burden weights retain recovery cost, while separate physical-loss and service constraints prevent a favorable composite from authorizing an unacceptable component outcome.

Moreau, L., & Missier, P., eds. (2013). PROV-DM: The PROV Data Model. W3C Recommendation. Distinct entities, activities and derivations inform the source-version and action-event crosswalk. Provenance makes changes auditable; it does not authenticate unavailable originals or identify causal benefit.

Private research memo records access scope: full text, primary abstract or continuing reference verified earlier in the series. No external effect size is transferred.

FINAL REFEREE REPORT / INTERNAL AGENT-ASSISTED REVIEW

Approve the design case within its source and identification boundaries

[AI] Separate evidence, methods and numerical reviews inform this author-led report. It is internal agent-assisted review, not external peer review, client approval or a certification of consulting-firm rank.

Major issue	Disposition in the revised case
Source identity and arithmetic	Both headline versions remain explicit; category nonclosure is preserved. Small displayed target discrepancies are evaluated against rounding intervals rather than automatically declared errors.
Benefit overlap and portfolio scope	The protected opportunity range isolates one pair under disjoint-other-action assumptions. It is not a universal portfolio floor. Projected entitlement is separated from observed spending and marginal returns.
Outcome and attribution	Physical nonreceipt, system classification and concessions remain different constructs. Another team’s reported technical movement does not become this author’s impact.
Control and evaluation design	The proposed field test retains original exception episodes, costly recovery and unknown truth. Ordinary-flow and physical-loss guardrails require separate precision.
Economics and reproducibility	Hypothetical burden reduction, incremental cost and information value use explicit models and independent numerical checks. Source integrity, privacy and actual-page inspections have separate records.

Partner critique: can a COO act on page one?

[AI] Yes: reconcile the distinct opportunity and authorize calibrated testing before making an additive benefit commitment. The sponsor owns rollout; operating owners demonstrate activation; independent evaluators certify the outcome evidence against declared thresholds.

Minor revisions and remaining limits

[AI] The cover states its overlap premise; benefit is compared with cost before netting, and source-derived role support is graded. Originals, historical implementation and mature effects remain unestablished. Field recruitment, oracle quality, penalty weights and additional service harms require calibration.

REVIEW DISPOSITION — PASS WITHIN STATED SCOPE

[AI] The case supports the documented diagnostic and operating proposal, new reconciliation findings and a reproducible prospective test. It establishes no realized program saving, causal customer benefit or deployment of the proposed controls.

Labels: [HR] historical record; [RC] recalculated; [AI] analytical inference; [PD] proposed design. Source basis S1; see private annex.