

OPERATING CASE STUDY / OPERATIONAL EXCELLENCE

Recovery Planning & Commitment Assurance

A green date needs a defensible recovery contract

An anonymized project case study of a service-network recovery plan: reconciling competing targets, testing station-level feasibility and connecting executive commitments to accountable operating actions.

PROJECT CONTEXT / DISTRIBUTED SERVICE OPERATIONS

BLUF / BOTTOM LINE UP FRONT

Validate the target.
Prove the route to recovery.

RECALCULATED FINDING

Recalculation finds an 8.41% difference between two thresholds for the same milestone. In a conditional station solve, a hypothetical 5% deterioration elsewhere raises the selected cohort's required reduction from 24.02% to 34.04%.

CONTRIBUTION IN THE SOURCE RECORD

Recovery planning, analytical challenge and target governance. Registry-attested responsibility; provisional hoop-level support.

EVIDENCE STATUS

Planning and evidence-reconciliation case. Target achievement, causal customer impact and realized savings remain unestablished.

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

EXECUTIVE ABSTRACT

Authorize recovery against one target contract and an evidenced action package

A service-network plan translated monthly goals into weekly accountability. Its record shows why the target, forecast and operating commitment need separate approval.

CONTEXT | A recovery promise crossed several planning systems

[HR] The register describes a post-period replan, selected-station analysis, weekly target tables, named owners and escalation triggers across two operating groups. Source artifacts use different target constructions and cumulative planning versions.

OBJECTIVE | Establish what a green date authorizes

[PD] Distinguish the approved performance threshold, the forecast of reaching it and the funded actions intended to deliver it. A scheduled milestone becomes a credible commitment only when all three refer to the same outcome and population.

APPROACH | Reconcile identity, feasibility and coverage

[RC] Compare target constructions and reconstruct the station-tier solve. [AI] Challenge rate aggregation, cumulative recovery timing and input-to-outcome coverage. [PD] Formalize a Recovery Commitment Contract, with separate authority for the target, forecast and operating package.

FINDINGS | Equivalent labels conceal different recovery burdens

[RC] Two thresholds for the same milestone differ by 8.41%. The selected cohort requires a 24.02% reduction under a frozen-rest assumption; a hypothetical 5% deterioration elsewhere raises that requirement to 34.04%. [AI] Neither a permissive target nor a named input owner establishes customer impact.

DECISION VALUE | Resolve the contract before committing the promise

[PD] The COO should name a commitment owner, ratify the target basis and require a reconciled forecast plus a feasible, resourced action package. Renew approval when scope, weights, target vintage or delivery assumptions change.

IMPLICATIONS | Credit planning within its evidence boundary

[AI] Registered ownership and co-preparer records provide provisional hoop-level support for planning and analytical contribution. They do not independently authenticate sole authorship, adoption, field execution or savings. No observed recovery milestone is established here.

S1: selected project register, M01–M33 and linked evidence. HR historical record; RC recalculated; AI analytical inference; PD proposed design. Original artifacts were not independently authenticated in this case.

EXHIBIT 1 / TARGET IDENTITY

Two thresholds permit different performance at the same named milestone

[RC] The step endpoint is 8.41% higher than the monthly-held threshold assigned to the same milestone. A common date does not make them equivalent targets.

THE SAME MILESTONE HAS TWO REGISTERED THRESHOLDS

Monthly-held threshold = 100 • lower is better • plan comparison, not an outcome



That conclusion requires the same event, exposure, population and period.

Approval lineage and construct equivalence remain unresolved.

No final-milestone actual is established in this case.

[HR] One artifact carries a declining weekly step series. Another repeats a monthly plan value across its assigned weeks. Both name the same recovery milestone, but their endpoint thresholds differ. [AI] These are planning constructions from different artifacts and vintages; their common label does not establish common approval or achieved performance.

Contract field	Required decision
[PD] Outcome and scope	Confirm customer-event definition, eligible exposure, operating population, aggregation and maturity window.
[PD] Clock and authority	Identify the applicable period, target version, approving owner and effective supersession rule.

[AI] If the event and exposure contracts match, the higher threshold permits 8.41% more recorded defects per fixed exposure. If they do not match, the ratio is only a comparison of registered numbers. Neither interpretation supports treating the two as interchangeable. A weekly actual also cannot be judged against a cumulative goal merely because the population matches.

EXECUTIVE ACTION

[PD] Ask the target owner to ratify the governing construction before releasing a recovery promise. Preserve both versions and the reason for selecting one. Do not average the thresholds or silently replace an earlier decision record.

[HR/RC] S1:M01/M19/D03/D07; E01/E03/E08/E09. Monthly-held threshold indexed to 100. Same milestone is a registered plan label, not an observed result. Population equivalence and approval lineage require verification.

EXHIBIT 2 / CONDITIONAL FEASIBILITY

The cohort recovery burden rises when performance elsewhere deteriorates

[RC] The registered tier solve requires a 24.02% reduction from the selected cohort. A hypothetical 5% increase in the rest-of-network rate raises the requirement to 34.04%.

REST-OF-NETWORK DRIFT RAISES THE REQUIRED COHORT CUT

Group A • fixed exposure mix • reduction relative to the selected cohort baseline



Reconstructed cohort exposure share: 22.99%; it is not the site share.

These are conditional accounting scenarios, not forecasts or treatment effects.

[HR] The source solves for the cohort rate needed to reach the network target while holding other sites at their baseline. [RC] Its rounded inputs imply that the cohort carries 22.99% of shipment exposure. A fraction of sites and a fraction of shipments are different weights; the site-selection label cannot substitute for the exposure share.

Feasibility assumption	What must be verified
[PD] Fixed population and mix	Retain cohort membership, eligible exposure weights and comparable definitions across baseline and target periods.
[PD] Achievable action effect	Establish the delivered package, realistic effect range, capacity and cost. A solved target is not evidence that teams can achieve it.

[AI] The sensitivity changes only the rest-of-network rate; it is not an observed deterioration or a forecast. Selecting the worst baseline sites also creates regression-to-the-mean risk. Improvement in their subsequent rank or rate cannot establish treatment impact without a valid comparison.

COMMITMENT TEST

[PD] Require the operating owner to accept a scenario range, monitor exposure mix and reopen the commitment when performance outside the cohort exceeds the allowed envelope. Evaluate causal delivery separately from the accounting solve.

[RC] S1:M14/D01 and registered rounded tier inputs. Drift scenarios are hypothetical; all other inputs remain fixed. Required reductions use the cohort baseline. Implied exposure weights are reconstructed, not independently observed.


EXHIBIT 3 / THRESHOLD SLACK

A campaign can be unnecessary for the target and still valuable for customers

[RC] Group B’s future threshold is 6.40% above its frozen current rate. Holding that rate would satisfy this threshold without any improvement.

TARGET SLACK DOES NOT MEASURE CUSTOMER VALUE

Group B • frozen current network rate = 100 • future threshold on that same scale



A no-campaign decision answers the first question only under its stated assumptions.

[HR] The source recommends no selected-station campaign because the network already meets the future threshold under its fixed-baseline solve. [AI] This is a conditional allocation argument. It does not establish that the tail cannot become binding, that the network will remain below target or that correcting tail defects has no value.

Two separate questions	Decision criterion
[PD] Is the campaign necessary for this target?	Test the authorized threshold against plausible changes in the rest of the network, exposure mix and measurement.
[PD] Is the campaign worth undertaking?	Compare incremental customer harm avoided with delivery cost, scarce capacity and alternative uses; retain safety and service guardrails.

[AI] The future threshold is higher than the current rate; this is target slack, not a causal performance gain. The most extreme site rate also need not represent the greatest absolute customer burden. Volume, severity and feasible response matter alongside rate, with small-population uncertainty retained.

COO CHOICE

[PD] Accept a no-campaign recommendation only for its stated objective and assumptions. Assign an owner to monitor the conditions that would reopen it. A green aggregate status cannot waive an independently material customer or safety obligation.

[HR/RC] S1:M15/D02/E02/E10. Current network rate indexed to 100; future threshold =106.40. The 6.40% uses the current-rate denominator. Historical adoption and effects of the campaign recommendation are unestablished.


EXHIBIT 4 / COMMITMENT AGGREGATION

Monthly rate commitments need exposure weights before they become a benefit

[AI] Adding planned monthly reductions produces a sum of rate commitments. It does not establish unique customer events prevented or economic value.

MONTHLY COMMITMENTS ARE RATES BEFORE THEY ARE BENEFITS

Group A • planned monthly rate reduction • first month = 100 • no realized effects



$$\text{Weighted mean reduction} = \frac{\sum(\text{exposure} \times \text{rate reduction})}{\sum(\text{exposure})}$$

Prevented events also require incremental effects, aligned units and overlap control.
Actual monthly exposure weights were not recovered.

[HR] The source adds commitments across months and outcome categories to identify the largest recovery lever. [AI] Benefit validation requires each month’s eligible exposure, matched definitions and incremental action effects. A weighted average rate and a sum of monthly rates answer different questions.

Conversion	Evidence required
[PD] Rate to customer events	Multiply each validated incremental rate reduction by its own eligible exposure, then resolve overlapping events and interacting programs.
[PD] Events to economic value	Apply mature outcome-specific avoidable loss, reconcile credit elsewhere and deduct incremental delivery and validation cost.

[AI] Identical per-exposure commitments do not become inherently harder merely because one network is smaller. Lower volume changes absolute event benefit; difficulty depends on the starting burden, addressable share, capacity and cost. Nor does a large planned share prove that the selected lever caused the historical miss or will close the future gap.

FINANCE GATE

[PD] Keep plan reconciliation separate from benefit validation. Ask finance and operations to approve the exposure-weighted portfolio, its overlap rules and its delivery assumptions before translating a commitment share into an entitlement or saving.

[HR/AI] S1:M12/M13/M21/D01/E02/E04/E09. The displayed conversion is proposed, not a recovered customer benefit. Rounded monthly and category sums reconcile within reported precision; that does not independently validate the commitments.


EXHIBIT 5 / RECOVERY TIMING

A first recovery month needs a reconciled cumulative ledger

[HR/RC] One registered cumulative planning table stays positive throughout its horizon. Another version names a first favorable month. The record does not establish a common zero-crossing basis.

ONE PLAN VINTAGE SHOWS NO ZERO CROSSING IN EITHER GROUP

Recalculated cumulative plans • each group’s first listed balance = 100 • separate scales



A competing version names a recovery month; its balance and sign basis need reconciliation.
Positive is a mathematical sign here, not a verified favorable status.

[AI] A cumulative value is interpretable only with its opening balance, favorable-sign convention and monthly flow definition. If all displayed values have the same sign, that table cannot by itself identify a new zero crossing. It may begin favorable, remain unfavorable or answer a different question; the sign convention and scope determine which.

Recovery statement	Required basis
[PD] Current-period target met	Comparable actual and threshold for one eligible period; no automatic claim that earlier shortfalls were recovered.
[PD] Cumulative deficit recovered	Opening deficit plus matched monthly flows, a fixed favorable-sign convention and a documented crossing rule.

[HR] The source ties recovery timing to the specific plan vintage. [AI] Reforecasting a threshold, exposure or opening balance can move a projected date without an operational action. Neither a preferred narrative nor a later file timestamp establishes the governing construct or an achieved recovery.

VALIDATION HOLD

[PD] Require the model owner to reconcile the two ledgers and identify which decision each supports. Retain unresolved timing as unresolved. Publish a recovery date only with the approved basis, forecast uncertainty and evidence required to declare achievement.

[HR/RC/AI] S1:M17/M18/M19/M20/E02/E03/E04. Competing versions are planning records, not sequential performance observations. Favorable-sign and opening-balance comparability remain unresolved; no recovery achievement is established.


EXHIBIT 6 / MECHANISM AND OWNERSHIP COVERAGE

An input owner is not an outcome guarantee

[HR] A response draft acknowledges incomplete input-to-output coverage. [AI] This identifies missing mechanisms, not the causal amount needed to recover.

OWNERSHIP MUST CONNECT TO A VERIFIED CUSTOMER MECHANISM

Proposed evidence chain • each arrow requires its own population and verification



Retain uncovered mechanisms and overlapping program effects.

The historical record supports planning responsibility, not sole authorship or field results.

[HR] The record names input owners and identifies output categories needing additional mechanisms. [AI] An input indicator is not an additive slice of the customer outcome. An observed category share is also not the share of the target gap that an intervention can close. Those relationships require a delivered action package and a customer-outcome estimand.

Crosswalk layer	Evidence and responsibility
[PD] Input → action → outcome	Operating owners specify the change, exposure population, delivery fidelity, maturity window and plausible downstream mechanism.
[PD] Project → contribution	Separate the principal’s registered planning and analytical work from co-preparation, source-workbook ownership and field implementation.

[AI] Registered goal-owner fields and co-preparer metadata offer provisional hoop-level support for planning responsibility. A retrospective memory transcription offers weaker authentication of the target email. A draft acknowledgment does not establish that it was sent or that its promised follow-up was delivered. These boundaries preserve the value of the reasoning without expanding the historical claim.

EXECUTION GATE

[PD] Assign every material recovery assumption to an accountable action owner and a verification record. Keep uncovered mechanisms, overlapping effects and unresolved contribution claims visible. Named ownership is the start of assurance, not proof of impact.

[HR/AI] S1:M29/D04/D05/E01–E05/E07–E09/E12. Evidence grade: provisional hoop for registered planning responsibility; no smoking-gun or doubly decisive test passed. Field delivery, sole authorship and historical causal effects remain unestablished.


EXHIBIT 7 / RECOVERY COMMITMENT CONTRACT

Approve the target, forecast and action package separately

[PD] The Recovery Commitment Contract joins one outcome definition to separate target authority, predictive evidence and operating responsibility.

ONE RECOVERY CONTRACT REQUIRES THREE DISTINCT APPROVALS

Proposed framework • one commitment owner reconciles the approvals



Approval gate	Accountable evidence
[PD] Target authority	Target owner approves scope, clock, version, threshold and achievement rule; records supersession.
[PD] Forecast validity	Analytical owner supplies as-of forecasts, uncertainty and out-of-sample validation. A target ramp is not a forecast.
[PD] Operating commitment	Delivery owner accepts funding, effect range, dependencies, capacity, coverage and verification.

[PD] The COO appoints one commitment owner to reconcile approvals and escalate conflicts. Reopen after material changes in target, population, exposure weights, measurement or delivery assumptions. Preserve earlier approvals so a changed contract does not masquerade as improvement.

COO DECISION

[PD] Authorize reconciliation and bounded evaluation. Test prediction separately from causal delivery. Release the promise when target, forecast and delivery approvals agree. Keep planning contribution visible while withholding unestablished achievement, customer-impact and financial claims.


[PD] Luna Sol proposed Recovery Commitment Contract, informed by S1:D01–D08. Forecast and prospective evaluation methods are specified in the technical layer. This framework is not an authenticated historical deployment or an external certification.

A recovery contract influences action before it changes risk

[PD] The Recovery Commitment Contract separates an agreed target, a conditional forecast, an authorized action and a measured customer outcome.

CAUSAL DELIVERY IS DISTINCT FROM TARGET STATUS

Proposed DAG • no historical edge effect is estimated • confounder edges listed below



C: demand, mix, staffing, refund policy and concurrent programs → A, X and Y.

Primary trial path: Z → A → X → Y; retain all upstream-eligible units, including holds.

Measurement paths: Y → M; Q → M; L → M. Compare M with T for status only.

Outcome clock: original scheduled dispatch + 56 days. Later reviews belong to future cycles.

Unexplained residuals permit unmeasured paths; they do not identify a cause.

[PD] P is the recovery contract; D the resource decision; A delivered pre-dispatch action; X operating state; Y the fixed-horizon eligible-unit concession; M its measurement. Randomized assignment Z encourages A. Target identity T and forecast F inform D. C includes demand, mix, staffing, refund policy and concurrent programs. Quality Q and reporting lag L affect M.

[AI] Registered planning artifacts support P and its proposed D→A pathway. The tier equation is an accounting constraint, not an estimate of A→Y. Input correlations concern X and M; they do not identify the package effect. Historical customer and financial effects remain unestimated.

[PD] Compare reported M with the applicable T to describe status; do not treat that comparison as evidence of causality. The prospective trial identifies assignment Z→Y without conditioning on adherence A or mediator X. Index repeated reviews in time to avoid a same-period feedback cycle. Unexplained residuals are compatible with omitted paths, not proof of any named cause.

Source S1: registered plan, action dependencies and input/output distinctions. The DAG is proposed; no historical path effect is identified.

TECHNICAL FOUNDATION · ESTIMANDS AND STATUS

Specify the target event before assigning a green date

[PD] A target contract fixes scope, outcome, denominator, clock, period, horizon, authority and vintage. The same number can represent different objects.

Quantity	Definition and inferential boundary
Target / status	[PD] Desired rate g and observed compatible rate y . For lower-is-better outcomes, slack $s=g-y$. Its change is $\Delta s=\Delta g-\Delta y$: target movement and performance movement remain separate.
Forecast	[PD] Distribution of a defined future Y given information available at issue time and an explicit action scenario. A target path is not that distribution.
Conditional tier requirement	[RC] Tail rate needed for the network to meet g while non-tail performance and exposure mix remain fixed. This is a constraint, not an achievable causal effect.
Package policy ITT	[PD] Difference in 56-day risk of at least one prespecified all-cause service concession per upstream-eligible customer unit under assigned package versus control, averaged over its first four weeks.
Delivered-action effect	[PD] Effect of receiving the package. Nonadherence means this differs from assignment ITT; an as-treated comparison needs additional identification.

[PD] Freeze eligibility and site before checks; anchor 56 days to original scheduled dispatch. Retain held, cancelled and nonadherent units; bound unresolved missing outcomes. Benefit index= $100\times$ absolute risk reduction/reference risk, using one pretrial reference, not current control risk. Equal exposure weights and common bin-specific effects underpin power; changed weights require recalculation.

[AI] Weekly rate compliance, monthly rate compliance and cumulative-unit recovery are distinct events. Neither a projected crossing nor a rewritten goal supplies an observed policy effect.

Source S1: multiple target meanings, periods and projected green dates. Hyndman–Athanasopoulos distinguishes forecasts, goals and plans. Trial estimand is proposed.

TECHNICAL FOUNDATION • MEASUREMENT VALIDITY

A common label does not establish a common measure

[HR] The register contains competing target vintages, input definitions and population mappings. These are measurement conditions to resolve before estimating effects.

Metric family	Validity contract
Concessions	[PD] Distinguish unique shipments with a concession from recurrent concession events. Freeze shipment eligibility, origin site, policy, cohort clock and adjudication horizon.
Late / lost / other causes	[PD] Test whether causes partition the same outcome universe. Preserve unidentified residuals and overlapping pathways; taxonomy changes need a crosswalk.
Sort and resolution	[PD] First-pass success, same-day compliance, resolution effectiveness and tool adoption use different constructs or denominators. Do not substitute one to fill another's history.
Missing-dispatch indicator	[PD] Match dispatched exposure and event definition; reporting corrections and reclassification can change measured levels without changing customer outcomes.
Target and commitment	[PD] Record authority, version and scope. An owner named beside a goal is not proof that the action was approved, delivered or effective.

[PD] Proposed validation: independently double-code 300 records across metric, node, period and disputed category. Blind coders to status and proposed action; weight oversampled disputes when pooling. Report confusion matrices, agreement and Cohen's kappa with cluster-bootstrap intervals; adjudicate after coding. Pilot prevalence and dependence before confirmatory sizing. Kappa alone is not construct validity.

[PD] Freeze issue-time and harmonized snapshots; reconcile exposure, revisions and delay. Test invariance across business lines. The proposed 56-day endpoint starts at original scheduled dispatch, including held/cancelled units. It is a future fixed horizon, not verified ultimate maturity. Missing outcomes require bounds, not zero coding.

Source S1: construct changes, conflicting targets and scope limitations. Wang–Strong supports fitness for use; Croushore–Stark supports preservation of data vintages.

TECHNICAL FOUNDATION • PORTFOLIO COHERENCE

Add avoided units after aligning exposure and joint treatment

[AI] Internal rollup agreement validates an accounting identity. It does not validate a six-month rate, portfolio feasibility or causal additivity.

[PD] Let d_m be an attributable reduction per million eligible shipments and V_m monthly eligible exposure. $\text{Avoided units} = \sum V_m d_m / 10^6$; $\text{period rate reduction} = \sum V_m d_m / \sum V_m$. Summing monthly rates produces rate-months. Equal-month exposure is a scenario, not an established property of the source.

Coherence test	Decision requirement
Time aggregation	[PD] Reconcile weekly targets to calendar months using shipment weights and weeks crossing month boundaries. Flat allocation adds accountability detail, not predictive information.
Portfolio inclusion	[HR] A ramping subset sits within the broader commitment. Count it once; enumerate packages, populations, start dates and maturity assumptions.
Joint effects	[PD] For two packages, $\delta AB = E[Y_{00} - Y_{11}]$. Y_{00}/Y_{11} denote neither/both packages; Y_{10}/Y_{01} denote each alone. It need not equal $E[Y_{00} - Y_{10}] + E[Y_{00} - Y_{01}]$. Shared beneficiaries and interaction require joint-outcome evidence.
Economic translation	[PD] Apply avoidable cost and nonduplicated value to the same mature-outcome estimand, then subtract incremental cost. Do not add the plan total to component-program benefits.

[AI] Equal DPMO commitments in differently sized networks imply equal rate reductions, not equal unit reductions or automatically greater difficulty in the smaller network. Difficulty depends on baseline risk, actionable mechanisms, capacity and cost.

[PD] Recover exposure forecasts and joint eligibility before treating a portfolio bridge as feasible. A factorial design could isolate interactions if operationally feasible; the proposed package trial estimates the combined policy instead.

Source S1: monthly commitments and nested ramp assumptions. Athanasopoulos et al. supports temporal coherence; VanderWeele–Knol establishes scale-dependent interaction.

TECHNICAL FOUNDATION · SELECTION AND FEASIBILITY

The tier requirement changes when the rest of the network moves

[RC] The selected-tail solve is conditional on exposure share, non-tail risk and target authority. Site share is not necessarily shipment share.

[RC] With tail exposure share w , tail rate t and non-tail rate r , network rate $g=w t+(1-w)r$. Required tail rate $t^*=[g-(1-w)r]/w$. The reconstructed wider-tail share is about 23%; its required reduction is about 24.0%, assuming the recorded inputs refer to one population.

[RC] At that share, $\partial t^*/\partial r=-3.35$: one common-scale rate point of non-tail worsening requires 3.35 additional tail points of reduction. This sensitivity uses rounded inputs. Transport to the trial requires validated unit, denominator, maturity and weights; the shipped-exposure solve is not automatically its estimand.

Threat or moderator	Implication
Regression to the mean	[PD] Under a stationary bivariate-normal score model, correlation ρ retains ρ of selected baseline excess. Assumed $\rho=.25/.50/.75$ implies 75%/50%/25% expected excess reversion without treatment.
Scope and capacity	[PD] Prespecify node type, business line, baseline risk and intervention authority as moderators. A package cannot act on mechanisms outside its defined operating boundary.
Nonbinding target	[AI] Existing headroom can remove the need for a campaign solely to meet that frozen target. It does not establish zero customer or economic value from improvement.

[PD] Freeze selected membership and source vintage; compare contemporaneous eligible non-focus sites when available. Historical reliability and focus/non-focus effects are not estimable from the register. Randomize within the eligible population and retain sites after re-ranking or graduation.

Source S1: conditional tier solves, materiality ranking and different network decisions. Regression-to-the-mean values are explicit synthetic sensitivities, not observed adjustments.

TECHNICAL FOUNDATION • VARIATION AND SIGNAL

Engineered target steps are not observations of process stability

[RC] One short missing-dispatch series crosses tentative limits. The other comparable short sequences do not; none identifies a package effect.

[RC] $MR_t = |x_t - x_{(t-1)}|$; individual limits = $\text{mean}(x) \pm 2.66 \text{ mean}(MR)$; MR upper = $3.268 \text{ mean}(MR)$. Each sequence's first observation = 100. Short retrospective limits cannot certify stability; the same-phase assumption remains provisional.

Actual sequence	Tentative results on its own index
Group A missing-dispatch	[RC] Six observations; center 109.10; individual limits 92.03–126.16; MR upper 20.97. Final observation and final MR cross: investigate, without causal attribution.
Group B missing-dispatch	[RC] Six observations; center 104.19; individual limits 93.09–115.29; MR upper 13.64. No crossings; compatible with routine variation under these limits.
Group A resolution	[RC] Five observations; center 102.07; individual limits 94.73–109.41; MR upper 9.02. No crossings; compatible with routine variation under these limits.
Remaining ledger	[AI] Historical-version forks are sensitivities, not independent histories. Pairs, stitched scopes and unresolved hours-plan timing are non-diagnostic. Definition breaks prohibit causal narratives.

[PD] No binary count/denominator panel supports Laney p' here. Future use requires valid binary incidence and stable eligibility; recurrent events require a count model. Preserve revision forks and do not fabricate observations from ranges.

[AI] Engineered target steps are excluded from process-signal claims. Mechanical plan-shape checks are retained privately without a process-signal interpretation. Governance tripwires are management thresholds, not statistical falsifiers. Investigate reporting corrections, mix and shared shocks before assigning cause.

Source S1 and protected SPC ledger; NIST and Laney. Alternative source versions and exact series eligibility are retained privately. Public limits are indexed and rounded.

TECHNICAL FOUNDATION · FORECAST VERIFICATION

Test a frozen prediction against the event it actually names

[PD] Forecast reliability requires repeated issue-time predictions and compatible outcomes. A single projected green date cannot establish calibration.

Verification layer	Prespecified rule
Forecast case	[PD] Freeze scope, target vintage, issue time, horizon, action scenario and outcome maturity. Score competing methods on exactly the same cases.
Probability of green	[PD] Define $G=1\{Y\leq g\}$. For predicted probability p , use Brier loss $(p-G)^2$. Do not replace g after the outcome is known.
Full distribution	[PD] Use a proper score such as CRPS, alongside interval coverage and width. Narrow intervals alone do not demonstrate accuracy.
Reliability by use	[PD] Assess calibration by horizon, operating scope and intervention scenario; account for clustered sites and overlapping forecasts in uncertainty estimates.
Decision value	[PD] Link forecasts to a declared action set and payoff model. A better forecast score is not automatically a financial benefit.

[AI] Forecasts conditional on one action policy cannot be interpreted as no-action forecasts after management changes that policy. Revisions should preserve both as-issued predictions and newly informed predictions; hindsight updates must not overwrite the original scoring record.

[PD] Benchmark against a frozen seasonal or persistence forecast appropriate to the available history. Choose the comparison before viewing outcomes. Use holdout periods and report forecast errors separately from the causal trial: predictive accuracy neither proves nor requires that the forecasted input causes the outcome.

RELIABILITY GATE

[AI] Reconciled targets are a coherent contract. Calibration is an empirical property of predictions and realized outcomes; the case does not establish it historically.

Gneiting–Raftery (2007): proper scoring; Gneiting–Balabdaoui–Raftery (2007): calibration and sharpness. The verification protocol is proposed.

TECHNICAL FOUNDATION • PROSPECTIVE EVALUATION

Estimate the ramp instead of assuming instant full benefit

[PD] Randomize a specified pre-dispatch package; retain all upstream-eligible units. Its policy effect is separate from the effect of writing a recovery plan.

[PD] Eight enrollment weeks; balanced site-cluster sequences start in weeks 3, 5 or remain waitlisted through final ascertainment. Stratify by node, business line and baseline risk. Freeze eligibility before checks; anchor follow-up to original scheduled dispatch. Later rollout cannot retroactively treat earlier control cohorts.

Design requirement	Specification
Effect curve	[PD] Fit site/calendar effects plus first-two, next-two and later exposure bins. Primary benefit is minus the equal mean of the first two coefficients; assume common bin effects across sites/calendar.
Outcome and guardrails	[PD] Unique eligible-unit all-cause service-concession risk at 56 days, including held/cancelled units. Bound missing outcomes. Monitor holds, cancellation, unfulfilled demand, safety, severity and resources; Late-specific risk diagnoses mechanism.
Power and inference	[PD] Contrast-specific GLS under synthetic covariance; compare MDE and economic power with break-even. Required sites are not verified availability. Any variance-reduction redesign needs pilot evidence and finite-cluster calibration.
Contamination / fallback	[PD] Confine treatment to initial pre-dispatch actions. If routing or downstream treatment spills across cohorts or sites, use independent parallel clusters and recompute power.

[PD] Inspect baseline balance and event-time diagnostics; non-significant pretrends do not establish comparability. If randomization is impossible, a justified Callaway–Sant’Anna group-time design requires overlap, no anticipation and conditional parallel trends, with cohort-specific effects and clustered inference. It cannot identify effects from this register alone.

Proposed trial; Kenny et al. (2022) motivates exposure-time effects. Hemming et al. (2018) informs rollout reporting. Power/economic assumptions appear in the decision model.

TECHNICAL FOUNDATION • ATTRIBUTION AND LIMITATIONS

Credit plan architecture without borrowing field results

[HR] The register identifies the principal as a plan author and goal owner. It also records separate ownership of the underlying workbook, executive commitment, programs and field delivery.

Personal-contribution claim	Process-tracing grade
Authored recovery architecture	[HR] Provisional hoop support: registered authorship, decision box, accountability matrix and target audit align with the role. Original artifacts remain unauthenticated in this pass.
Created weekly accountability	[HR] Provisional hoop support: recorded target handoff and downstream reference use support involvement. They do not prove all sites adopted or obeyed it.
Corrected interpretation	[HR] Provisional hoop support for detecting incompatible target identities and stating incomplete mechanism coverage. Verified approval/change receipts are needed to establish implemented corrections.
Delivered outcome or cash impact	[AI] Unestablished. Source observations and program projections do not isolate this plan's effect or establish realized savings.

[PD] A hoop test checks a necessary condition; smoking-gun evidence must be reliable and unlikely under alternatives; doubly decisive evidence must also exclude those alternatives. Repeated register summaries from one lineage are not independent corroboration.

[AI] Consolidated limits: secondary evidence; unresolved goal/actual vintages; missing exposure and treatment panels; incomplete action receipts; unobserved maturity and severity; selected cohorts; shared programs and policy changes. Forecast calibration, adoption causality and customer/cash effects are therefore unestimated. Conditional algebra and synthetic prospective models remain reproducible within their stated assumptions.

Source S1: explicit ownership boundaries and missing originals. Process-tracing grades describe the evidence available to this case, not the principal's full work history.

TECHNICAL FOUNDATION • CONTRIBUTION AND REPRODUCTION

The Recovery Commitment Contract makes the plan testable

[PD] RCC binds target identity, an exposure-weighted action portfolio, owner/resource authorization and prospective verification. Its propositions are testable; this case does not establish their efficacy.

Proposition	Falsifiable implication / test
P1 • Identity before status	[PD] Compatible target contracts reduce erroneous authorization relative to untyped green dates. Test on fixed eligible decisions, including holds and deferrals.
P2 • Conditional feasibility	[PD] A frozen-rest tier solve loses its guarantee when non-tail drift or mix exceeds slack. Check the equation against prospective fixed-scope exposures and outcomes.
P3 • Coherence before benefit	[PD] Exposure-weighted, overlap-audited commitments predict aggregate effects better than summed monthly rates. Compare frozen same-case predictions using a proper score.
P4 • Outcomes before scale	[PD] Input compliance can coexist with an uneconomic package. Scale only when the benefit lower bound exceeds break-even and guardrails pass; an upper bound below break-even contradicts sufficiency. A spanning interval is inconclusive.

[AI] RCC applies Ittner–Larcker's linkage concern and Campbell/Goodhart's target-response problem. Shewhart separates variation from signal; Tucker–Edmondson distinguishes solving a local problem from changing its organizational causes. Reconciliation and proper scoring separate coherence from verification. The contribution extends recovery-plan design with explicit decision contracts; this descriptive case neither confirms nor contradicts those general theories.

[PD] The private appendix preserves cell provenance, source hashes, compatible version forks, raw recalculations, synthetic design/covariance, payoff tables and independent checks. Public analyses use relative changes or positive common-baseline indices. Reproduce from the protected inputs; do not reverse-engineer source precision from rounded public displays.

Proposed RCC synthesis; substantive links to measurement, target response, process control and forecast evaluation. Protected reproducibility package and final referee report accompany the case.

MODEL 1 / TREATMENT, MATURITY AND RAMP

Evaluate the delivered package across its first four exposure weeks

[PD] These four pages specify a future test and decision model. They do not estimate historical customer impact, realized savings or forecast accuracy.

[PD] Randomize independent eligible sites equally to launch in enrollment week 3, week 5, or a waitlist held through final follow-up. Enroll eight weekly cohorts of upstream eligible customer units before pre-dispatch checks. Retain held, cancelled and noncompliant units; a release check must not improve the denominator by excluding failures.

Design object	Definition and boundary
Action package	[PD] Prespecified pre-dispatch readiness controls: authorized staffing, exception resolution and release checks. Implementation finishes before initial departure. It does not include later on-road changes or retrospective treatment of old shipments.
Primary outcome	[PD] Unique eligible unit with at least one prespecified all-cause service concession within 56 days of its original scheduled dispatch date. Validate cohort linkage and maturity; keep unresolved ascertainment in missing-data bounds.
Primary causal estimand	[PD] Upstream-unit-weighted ITT risk reduction averaged equally over the first two and next two exposure weeks. A later-exposure effect is separate. Never condition on departure, adoption, completion or recovery status.
Guardrails	[PD] Track holds, cancellations, unfulfilled promises, lateness, cost/severity, safety, workload and displacement. Suppressed departures or unresolved facts cannot certify success.

[PD] Divide absolute risk differences by one common pretrial reference risk, indexed to 100. A 12-point reduction is 12% of that fixed reference, not a 12% relative decline from contemporaneous controls. Equal eligible volumes are assumed; changed weights or baseline uncertainty require new power.

[PD] If shared staff, returned shipments or downstream actions transmit treatment to control cohorts, redesign as a parallel cluster trial. The current power numbers do not apply automatically to that design.

Proposed site-randomized phased rollout; target population, treatment completion and maturity require pilot verification.

MODEL 2 / COMPATIBLE ECONOMICS AND TARGET SUFFICIENCY

A worthwhile intervention can still leave the recovery gap open

[PD] The economic test and the recovery commitment answer different questions. Both require the same eligible exposure and maturity horizon.

[PD] $Net = 100 \times r \times a \times q - K$. Here r is the first-four-week absolute ITT risk reduction divided by reference risk; 100 is the reference concession-cost pool for that eligible exposure and horizon. Avoidable/monetizable unit-cost share $a=.65$, nonoverlap $q=.80$ and incremental cost $K=3.5$ are assumed.

Assumed rate reduction	a = 40%	a = 65%	a = 90%
6%	-1.58	-0.38	+0.82
12%	+0.34	+2.74	+5.14
24%	+4.18	+8.98	+13.78

[RC] Break-even is $K/(100aq) = 6.73\%$. At an assumed 12% reduction, base net value is +2.74 normalized cost units. Neither the outcome effect nor the financial inputs are observed project results.

[RC/PD] Separate target illustration: the recorded frozen-rest solve requires a 24.02% cohort reduction for an 8% network gap. A hypothetical 12% reduction in THAT recorded metric covers 49.95% of the gap. This is not an automatic translation of the trial’s upstream-unit endpoint.

[PD] Transport to the target requires a validated event/denominator crosswalk, maturity, exposure weights and comparable populations. Constant cost per prevented event is also assumed. Revalue and re-power for severity changes; count setup, delivery and displaced work once over the first-four-week exposure. Set a/q to 1 only if already embedded in net cash.

TWO RELEASE GATES

[PD] Scale the package only when a compatible confidence interval clears economic break-even and harm guardrails pass. Promise the recovery milestone only when the complete, nonoverlapping portfolio covers the gap under credible mix and rest-of-network scenarios.

Synthetic cost sensitivity, conditional on an identified mature effect. The tier linkage uses source-derived ratios and frozen-rest assumptions.

MODEL 3 / EXPOSURE-TIME POWER

Power the ramp effect instead of pretending full impact arrives on launch

[PD] Fit site and enrollment-week effects plus three treatment-age indicators. The primary reduction is minus one half of each of the first two coefficients.

[PD] Covariance is $\sigma^2[\rho 11' + (1-\rho)AR1(\phi)]$, with site-period mean SD $\sigma=20$, persistent share $\rho=.30$ and serial $\phi=.40$. Assume Gaussian errors, independent sites, known covariance and common age-bin effects across sites/calendar. The curve is 8, 16, then 24 reference-risk points; its first-four-week mean is 12.

Sites / residual SD	Zero-null MDE	Economic power at 12
24 / 20	14.71	16.9%
72 / 20	8.49	41.2%
144 / 20	6.01	69.1%
189 / 20	5.24	80.4%
72 / 12	5.10	82.6%

[RC] With $c=(.5,.5,0)$, variance is $c'I^{-1}c$ after site/calendar adjustment. $MDE=(z.975+z.80)\times SE$. Economic power is $Pr(\text{lower } 95\% \text{ bound} > 6.73\%)$ at true mean 12%. The smallest balanced design meeting both $MDE \leq \text{break-even}$ and economic power $\geq 80\%$ has 189 sites; low/high SD 16/26 require 120/318.

[RC] Collapsing the ramp to one immediate-effect coefficient yields weights 0.890, 0.091, 0.019 on the three exposure bins in this design. It would average the assumed curve to 9.03%, rather than the declared 12% estimand. An estimator's convenience must not redefine the investment question.

[PD] The 72-site precision redesign passes only if a pilot validates residual SD12: a 64% variance reduction from SD20 under the same design. Baseline covariates or better measurement cannot be credited with that gain in advance. If unavailable, the test cannot settle this investment decision. Do not expand beyond the eligible population to obtain sample size.

[PD] Pilot weights, covariance, finite-cluster inference, uptake and maturity. Cluster shared-resource sites together when spillovers require it, then recalculate power.

Synthetic planning calculation; separate profiled and full-design information checks. Kenny et al. (2022) motivates modeling exposure-time effects.

MODEL 4 / DECISION VALUE OF INFORMATION

Verify the recovery contract before making the expensive commitment

[PD] This model values a diagnostic decision before deployment. It does not depend on a positive customer-outcome effect.

Action / dominant constraint	Supported	Scope invalid	Delivery shortfall
Commit fixed milestone	+12	-18	-14
Resolve and rebaseline	+5	+10	-5
Authorize staged evidence	+4	+2	+8

[PD] The mutually exclusive states are: scope and delivery supported; scope invalid regardless of delivery; and valid scope with a delivery shortfall. Priors are .45/.25/.30. Payoffs are invented decision-resource units, net of each action’s effort and delay. The best current action authorizes staged evidence.

[PD] A contract/feasibility audit returns the correct state with probability b and either wrong state with probability (1-b)/2. Choose the action with maximum posterior expected payoff. Accuracy and calibration must come from independently adjudicated pilot records; reviewer confidence is insufficient.

Assumed accuracy b	Gross EVSI	Net after audit cost 1
45%	0.000	-1.000
70%	1.663	0.663
90%	4.288	3.288

[RC] Current expected value is 4.70; perfect information yields 10.30, so EVPI is 5.60. Every EVSI lies between zero and that ceiling. Posterior actions and probabilities are reproduced in the private package.

[PD] Scenarios are sensitivities, not confidence intervals or realized decision savings. Replace assumptions before commissioning an audit. Do not add these values to customer savings if they price the same prevented loss.

Synthetic finite-state information model; exact action/state definitions, priors, likelihoods and cost are disclosed.

Read a recovery commitment as a claim that must be tested

[AI] The case’s contribution is a stronger authorization rule: a reconciled target, a validated forecast and an accepted action package answer different questions.

Claim class	Meaning in this case
[HR] Historical record	A statement or artifact described in the supplied project register. The underlying originals were not independently authenticated here.
[RC] Recalculated	Arithmetic from registered inputs or a declared synthetic model. The result retains its input class and assumptions.
[AI] Analytical inference	A reasoned interpretation or constraint. It does not establish a causal effect or an achieved milestone.
[PD] Proposed design	A future contract, evaluation, forecast protocol or planning scenario; not a claim of historical deployment.

[AI] The Recovery Commitment Contract separates target authority, predictive evidence and operating responsibility. The same milestone’s thresholds differ by 8.41%. In the conditional tier solve, hypothetical deterioration elsewhere raises the selected cohort’s required reduction. These calculations change the commitment question without establishing recovery.

[AI] Forecast literature distinguishes goals from predictions and requires prospective scoring. Process-control literature distinguishes routine variation from signals. Measurement and learning research explain why a named indicator or owner does not establish a changed operating cause. RCC applies these principles; this descriptive case does not experimentally validate the framework.

[PD] Reproduction chain: frozen register and source hash → selected cells and claim ledger → version forks and raw calculations → indexed exhibits → separately labeled synthetic model → independently checked PDF. The private package preserves source identities, formulas, design matrices, assumptions, review records and build files.

DISTRIBUTION

[PD] Share this anonymized PDF as the portfolio case. Keep its separate evidence annex and reproduction package private. They contain internal source locations, identities, exact dates and raw values.

S1: supplied project register. Public historical quantities are indexed or relative. Proposed study sizes, assumptions and information models are synthetic and explicitly labeled.

REFERENCES / PRIMARY METHODS AND SUBSTANTIVE LITERATURE

A target, a forecast and a treatment effect are different objects

Sources support the analysis and proposed design. They do not authenticate the project, validate its forecasts or establish its effects.

Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: Principles and Practice*, 3rd ed., §1.2, “Forecasting, goals and planning.” Official author text. Separates desired outcomes, predictive statements and plans.

Athanasopoulos, G., Hyndman, R. J., Kourentzes, N., & Petropoulos, F. (2017). Forecasting with temporal hierarchies. *European Journal of Operational Research*, 262, 60–74. Author manuscript. Temporal coherence; exposure-weighted rate identities here are case-specific applications.

Gneiting, T., & Raftery, A. E. (2007). Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102, 359–378. Author full text. Proper scores for frozen probabilistic predictions.

Gneiting, T., Balabdaoui, F., & Raftery, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society B*, 69, 243–268. Author full text. Reliability differs from narrow prediction intervals.

Kenny, A., Voldal, E. C., Xia, F., Heagerty, P. J., & Hughes, J. P. (2022). Analysis of stepped wedge cluster randomized trials in the presence of a time-varying treatment effect. *Statistics in Medicine*, 41, 4311–4339. Opened author preprint; exposure-time estimands.

VanderWeele, T. J., & Knol, M. J. (2014). A Tutorial on Interaction. *Epidemiologic Methods*, 3, 33–72. Author index and indexed manuscript verified; direct PDF unavailable. Additive and multiplicative effects require different assumptions.

Hemming, K., et al. (2018). Reporting of stepped wedge cluster randomised trials: extension of the CONSORT 2010 statement. *BMJ*, 363, k1614. Publisher-indexed guidance; direct access limited. Supports explicit sequences, exposure timing, correlation and endpoint reporting.

Croushore, D., & Stark, T. (2000). A Real-Time Data Set for Macroeconomists: Does Data Vintage Matter for Forecasting? Federal Reserve Bank of Philadelphia, WP00-6. Official text. Preserving as-issued data is a methodological analogy, not logistics evidence.

Original research, official guidance and author manuscripts. Access and interpretive limits are recorded in the private methods memorandum; no source supplies a historical effect for this project.

REFERENCES / PRIMARY METHODS AND SUBSTANTIVE LITERATURE

Measurement and learning come before performance credit

Sources support the analysis and proposed design. They do not authenticate the project, validate its forecasts or establish its effects.

Wang, R. Y., & Strong, D. M. (1996). Beyond Accuracy: What Data Quality Means to Data Consumers. *Journal of Management Information Systems*, 12(4), 5–33. Author-hosted text. Data quality includes contextual fitness for the decision.

OECD / European Commission JRC (2008). Handbook on Constructing Composite Indicators: Methodology and User Guide. Institutional guidance. Construct definition, normalization, weights and sensitivity matter when combining indicators or decision-loss components.

Ittner, C. D., Larcker, D. F., & Randall, T. (2003). Performance implications of strategic performance measurement in financial services firms. *Accounting, Organizations and Society*. Author institutional summary. Measurement choices require validated links to performance.

Campbell, D. T. (1976). Assessing the Impact of Planned Social Change. Original occasional-paper record, ERIC. Evaluative pressure can change indicator behavior; the case does not establish gaming.

Goodhart, C. A. E. (1984). Problems of Monetary Management: The UK Experience. In *Monetary Theory and Practice*, 91–121. Publisher bibliographic record only. Used for the target-response concern, with no quotation from unavailable full text.

Tucker, A. L., & Edmondson, A. C. (2003). Why Hospitals Don't Learn from Failures. *California Management Review*. Author working paper. Distinguishes resolving an immediate problem from changing its organizational causes.

NIST/SEMATECH. e-Handbook of Statistical Methods: Individuals Control Charts. Official guidance on successive moving ranges and individual limits. Short retrospective limits in this case are tentative; designed target ramps are excluded from process-signal claims.

Laney, D. B. (2002). Improved Control Charts for Attributes. *Quality Engineering*, 14(4), 531–537. Publisher abstract verified. Extra-binomial variation requires appropriate denominators and dispersion estimation; no unsupported Laney chart is fitted here.

Original research, official guidance and author manuscripts. Access and interpretive limits are recorded in the private methods memorandum; no source supplies a historical effect for this project.

REFERENCES / PRIMARY METHODS AND SUBSTANTIVE LITERATURE

Selection, attribution and value require compatible evidence

Sources support the analysis and proposed design. They do not authenticate the project, validate its forecasts or establish its effects.

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. Extreme baseline selection can create apparent subsequent recovery.

Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20, 37–46. Chance-corrected coder agreement. Kappa is neither an independent ground truth nor sufficient construct validity.

Collier, D. (2011). Understanding Process Tracing. *PS: Political Science & Politics*, 44, 823–830. Necessary and discriminating evidence for contribution claims. Repeated records from one lineage are not independent corroboration.

Hernán, M. A., & Robins, J. M. Causal Inference: What If. Author-hosted text, revised edition. Treatment definition, causal graphs, selection and policy estimands. Supports upstream eligibility and assignment-based analysis.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*. Reopened author manuscript. Group-time effects are an observational fallback only when its identification assumptions are justified.

Roth, J. (2022). Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends. *American Economic Review: Insights*, 4, 305–322. A nonsignificant pretrend test does not establish comparability or validate an observational design.

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*. Information value concerns an improved decision, not an assumed positive intervention effect.

The W. Edwards Deming Institute. The Funnel Experiment. Institutional account of tampering with a stable process. Supports separating routine variation from an action trigger; it does not certify historical process stability here.

Original research, official guidance and author manuscripts. Access and interpretive limits are recorded in the private methods memorandum; no source supplies a historical effect for this project.

FINAL REFEREE REPORT / INTERNAL, AGENT-ASSISTED REVIEW

The revised case supports credible commitments with bounded claims

[AI] The author and three parallel agents reviewed source fidelity, methods and executive use. This is internal review, not external academic or client certification.

Major issue	Resolution in the revised case
Source and contribution	Register lineages and conflicting versions remain explicit. Planning contribution receives provisional hoop support. Unauthenticated originals, co-preparation and separate program ownership prevent sole-delivery or savings claims.
Treatment-based selection	Eligibility is frozen upstream. Held, cancelled and noncompliant units remain; follow-up starts at original scheduled dispatch. The endpoint and guardrails prevent exclusion through release decisions.
Estimand and economics	Common age-bin effects and one fixed reference-risk scale are explicit. The first-four-week cost pool matches the causal contrast. Trial effects are not automatically translated into the historical target metric.
Precision and feasibility	Independent QR and exact-information calculations pass 56 checks. The precision redesign requires pilot-verified variance reduction; inadequate eligibility or precision remains an unresolved investment decision.

Minor revisions completed

[AI] Threshold slack uses an explicit denominator. Rate-months remain distinct from weighted rates and benefits. No cumulative date is selected by narrative preference. Public historical values remain relative or indexed; raw levels, identities and cross-case inversion risks are kept out of the public artifact.

Partner critique: would the COO act on page 1?

[AI] Yes: ratify the target and require a feasible recovery route. The 8.41% threshold fork and the conditional 24.02%→34.04% cohort burden change make the decision concrete. Neither finding claims achieved recovery. The COO can assign authority, request validation and withhold an unsupported commitment.

REVIEW DISPOSITION

[AI] Suitable as an anonymized planning and operating-method case after recorded artifact checks. Historical adoption, customer impact and cash results require additional evidence. The proposed design remains conditional on its empirical gates.

Internal author-led review and independent agent challenge. Source audit, numerical reproduction, artifact checks and full-page visual review are recorded in the private package.