

Performance Measurement & Credit Integrity

A repaired classification does not identify a programme's causal effect

An operating-project case on challenging benefit attribution after a delivery-control launch. It separates customer outcomes, claim access and reporting classifications, reconciles benefit claims, and proposes a controlled test of credit-review decisions.

MEASUREMENT INTEGRITY & PROGRAMME CREDIT

BLUF / BOTTOM LINE UP FRONT

Validate the measurement and the counterfactual before recognizing programme value.

RECALCULATED FINDING

Analytical finding: the same flat claim total and changed operating split can coexist with programme benefit, no effect or harm. A reporting diagnosis cannot, on its own, settle causal credit or exonerate an intervention.

CONTRIBUTION IN THE SOURCE RECORD

Provisional contribution: authored a benefit-attribution challenge and follow-up; programme delivery and the diagnostic ruling belonged to other owners.

EVIDENCE STATUS

Original artifacts were not recovered. The case supports evidence reconciliation and a proposed review test; historical causal impact and savings remain unverified.

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

STRUCTURED ABSTRACT / DECISION BRIEF

Repair the inference before assigning the credit

[AI] A broken classification feed changes what an operating metric means. It does not, by itself, identify whether a concurrent programme helped, harmed or had no effect.

Context

[HR] A delivery programme introduced risk-adaptive verification controls and reported customer, productivity and financial benefits. A regional operating executive challenged attribution of a contemporaneous metric movement. The metric owner subsequently described a failed claim-classification feed and a changed operational split while total recorded claims reportedly remained flat.

Objective

[HR] The executive sought evidence for the attribution and requested additional drivers. The supported personal contribution is a written challenge and follow-up, provisionally evidenced by register extracts. Programme design, experimentation, deployment and the diagnostic ruling belonged to other owners.

Approach

[AI] The Credit-and-Measurement Identification Contract (CMIC) separates physical receipt, access to legitimate claims, recorded classifications and programme credit. Source reconciliation, causal counterexamples and a prospective shadow-review design test what the evidence can support without manufacturing a historical impact estimate.

Findings

[AI] The same flat claim total and changed classification split can coexist with programme benefit, no effect or harm. Similar experiment baselines do not establish duplicated benefits. Financial components can be compatible with a cumulative figure at reported precision while the annualized valuation bridge remains unresolved.

Decision value

[PD] The metric owner should label the affected view as unsuitable for causal credit until its population, capture and classification paths are validated. Finance should require a common valuation basis. The review owner should test CMIC in shadow mode and scale only when its decision-resource gain clears full incremental cost and its false-credit safeguard passes.

Implications

[AI] Preserve the challenge as an evidence-governance contribution. Keep the reported ruling distinct from an identified causal null. The new evaluation concerns credit-review decisions; it does not certify the historical delivery programme, quantify customer benefit or establish realized savings.

[HR] Historical record in selected register S1; originals not recovered. [RC] Recalculation; [AI] analytical inference; [PD] proposed design. Historical quantities and identities remain private.


EXECUTIVE EXHIBIT 1 / MEASUREMENT CONTRACT

Keep customer truth separate from the recorded claim

[HR] The reported failure affected the channel feeding a claim classification. Customers could use alternative channels, but the record does not establish complete capture or a valid reconstructed split.

MEASUREMENT CONTRACT / FOUR DIFFERENT OBJECTS

A reported classification is the end of a measurement path, not direct customer truth



Layer	Question that must be answered
Physical receipt	Did the original customer receive the parcel within the prespecified outcome window? A reporting label is not a direct observation of this construct.
Legitimate access	Could an eligible customer complete a valid claim through an available route? Fewer recorded contacts can reflect less need, lower access or changed routing.
Classification	Which captured claims receive the operational label, from which source fields and under which version? A failed exclusion feed can inflate a derived category.

[AI] Let F denote physical failure, A claim access, R recorded claims and O the operational classification. Changes in F, A or classification rules can change O. If only the classification feed fails, restoring it can correct O without changing F. If access also changes, even total R can lose comparability.

EXECUTIVE ACTION

[PD] Require an all-channel reconciliation and an independent outcome-reference sample. Protect legitimate customer access as a separate objective; never count blocked claims or reduced reporting as service improvement.

[HR] Historical record in selected register S1; originals not recovered. [RC] Recalculation; [AI] analytical inference; [PD] proposed design. Historical quantities and identities remain private.

EXECUTIVE EXHIBIT 2 / NONIDENTIFICATION

A classification diagnosis does not exonerate a programme

[AI] A new source-independent construction shows why the recorded diagnosis and the programme’s causal effect answer different questions.

IDENTICAL OBSERVATIONS / DIFFERENT CAUSAL EFFECTS

Wholly hypothetical: 1,000 original deliveries per period; claim capture held complete

BENEFIT WORLD	NULL WORLD	HARM WORLD
Observed post claims: 100	Observed post claims: 100	Observed post claims: 100
Without programme: 120	Without programme: 100	Without programme: 80
Effect on claims: -20	Effect on claims: 0	Effect on claims: +20

All worlds: operational label 60 to 80; other label 40 to 20.

The same observed total and split do not identify the untreated post-period outcome.

[PD] In the hypothetical construction, each period contains 1,000 original deliveries and 100 observed claims. The operational category rises from 60 to 80 because 20 exclusion records disappear. The post-period untreated counterfactual can still be 120, 100 or 80 claims. Those worlds imply a programme effect of -20, 0 or +20 claims, with identical observed totals and classifications.

[AI] In this construction, claim capture is complete and held fixed, so the ambiguity already exists before adding access failures. A separate capture construction shows that a flat observed total can also conceal a changed physical-failure burden. Neither construction estimates the historical effect.

What the ruling establishes

[HR] The metric owner recorded an alternative explanation and said the programme was not the contributor. [AI] That establishes the stated interpretation and a plausible measurement mechanism. It does not supply the missing untreated outcome, prove the incident had no effect or show that the programme could not have an offsetting effect.

DECISION THAT CHANGES

[PD] Reopen causal credit separately from metric repair. Authenticate the feed diagnosis and recovery interval, then obtain a valid programme comparison. A corrected dashboard is necessary measurement work, not a counterfactual.

[HR] Historical record in selected register S1; originals not recovered. [RC] Recalculation; [AI] analytical inference; [PD] proposed design. Historical quantities and identities remain private.


EXECUTIVE EXHIBIT 3 / POPULATION AND TREATMENT

Transport an experiment only after its treatment and population match

[HR] The register contains programme-reported experiments, selected high-risk cohorts and a regional significance claim. Original assignment records, effect estimates and analysis files were not recovered.

TRANSPORT CONTRACT / DO NOT SKIP THE BRIDGES

A programme-reported experiment needs authentication before extension



Required bridge	Why the missing bridge matters
Treatment version	A hard verification block and a soft default are different interventions. Labels alone cannot establish that a regional experiment tested the shipped package.
Target population	Selected high-risk stops, addresses, transporters and the whole delivery network have different weighting and exposure. A large selected-cohort decline is not a network effect.
Outcome and analysis	A significance threshold supplies neither effect magnitude nor independent replication. Assignment, clustering, claims maturity, missingness and multiplicity remain material.

[AI] Selecting extreme performers creates regression-to-the-mean risk in an uncontrolled before/after comparison. A valid randomized comparison within the selected population can address that bias, but does not automatically identify an effect in other populations. The reported experiment cannot be declared invalid merely because the targeted group began at high risk.

REQUIRED EVIDENCE

[PD] Recover assignment and treatment versions before transport. Predefine the original cohort and maturity window. Do not adjust for realized offline behaviour, delivery success or contact channel if the treatment can change those variables.

[HR] Historical record in selected register S1; originals not recovered. [RC] Recalculation; [AI] analytical inference; [PD] proposed design. Historical quantities and identities remain private.


EXECUTIVE EXHIBIT 4 / ECONOMIC RECONCILIATION

Resolve the valuation basis before alleging double counting

[RC] The component sum can match the reported cumulative total within stated rounding precision. The relationship to the annualized headline remains a basis question requiring evidence.

VALUE CONTRACT / A SUM IS NOT A VALUATION BRIDGE

Reconcile treatment, population, period and resource basis before recognizing value



Reconciliation test	Defensible conclusion
Time and valuation	An annualized run-rate, prior-period benefit and cumulative total need a common horizon and recognition basis. Arithmetic compatibility is not a verified financial bridge.
Geography and exposure	Worldwide claims cannot be divided by a regional cost pool without a documented allocation. A shared reporting node establishes a reconciliation need, not the amount of duplicated value.
Interactions and capacity	Overlapping mechanisms may interact. Near-identical baselines alone do not prove nonadditivity. Contact hours and route minutes require paid-capacity and overlap evidence before cash conversion.

[AI] The separate factorial construction holds each component’s standalone contrast fixed while its combined outcome varies across additive, redundant and adverse-interaction worlds. Joint assignment or defensible interaction assumptions are needed to value a package. A common baseline neither proves nor rules out incremental benefit.

FINANCE ACTION

[PD] Give every claim a treatment version, eligible population, valuation horizon, causal contrast, avoidable-cost share and overlap rule. Preserve unresolved claims as unresolved; neither book their total nor erase them by assertion.

[HR] Historical record in selected register S1; originals not recovered. [RC] Recalculation; [AI] analytical inference; [PD] proposed design. Historical quantities and identities remain private.


EXECUTIVE EXHIBIT 5 / CONTRIBUTION AND CLOSURE

Credit the written challenge without inventing its downstream impact

[HR] The register reproduces a written attribution question and a later follow-up. It also records a metric-owner reply and management questions about onset, repair and data recovery.

CONTRIBUTION CONTRACT / KEEP DECISION RIGHTS VISIBLE

Registry-level authorship is provisional; downstream impact needs separate evidence



Contribution claim	Process-tracing disposition
Question and follow-up	Provisional hoop support at registry level: quoted messages identify the executive as author. Original authentication remains open.
Detailed analytical standard	The richer maturity and segmentation protocol appears in narrative reconstruction. Do not elevate it to authenticated verbatim authorship without the original record.
Programme and root cause	No supported personal ownership of programme engineering, experimental execution or the diagnostic ruling. Shared visibility and copied recipients do not establish contribution.

[AI] A reply following escalation does not prove escalation caused the reply or shortened resolution. The metric owner had already received the earlier challenge. Missing repair and recovery dates establish incomplete documentary closure; they do not prove permanent uncorrectability or that no investigation occurred elsewhere.

CLOSURE CONTRACT

[PD] Assign the metric owner to the affected-cohort label, the platform owner to source recovery, the programme owner to causal evidence and Finance to valuation. Reopen on authenticated new evidence, with explicit due dates and decision authority.

[HR] Historical record in selected register S1; originals not recovered. [RC] Recalculation; [AI] analytical inference; [PD] proposed design. Historical quantities and identities remain private.


EXECUTIVE EXHIBIT 6 / PROSPECTIVE DECISION TEST

Test whether the review changes decisions at acceptable cost

[PD] A shadow evaluation fits the supported reviewer contribution. It compares a prespecified credit-assurance package with usual review without changing live financial authorization or customer treatment.

SHADOW TEST / EVALUATE CREDIT DECISIONS WITHOUT LIVE AUTHORIZATION

Randomize independent review domains; preserve every assigned docket



Design element	Prespecified requirement
Independent domain	Randomize review domains that share neither analysts nor decision artifacts across arms. Retain all assigned dockets, including deferrals, incomplete evidence and difficult cases.
Reference and policy	Blind the mature reference assessment to review arm. Freeze signal likelihoods, decision rules and loss weights before the evaluation; validate the reference independently.
Decision outcomes	Estimate decision-resource gain and false credit per assigned docket. Fully charge verification, delay and mistakes; do not equate a review gain with prevented physical loss.

[AI] The hypothetical finite model in the technical layer makes the operating law and cost assumptions explicit. Its power analysis tests economic detectability under selected allocations and stress conditions. Conditional model performance is a planning result, not evidence that CMIC has already improved historical decisions.

SCALE RULE

[PD] Require a lower confidence bound on decision-resource gain above incremental cost and a lower confidence bound on the reduction in false credit above zero. If the experiment cannot detect break-even or the reference is unreliable, revise the test before expansion.

[HR] Historical record in selected register S1; originals not recovered. [RC] Recalculation; [AI] analytical inference; [PD] proposed design. Historical quantities and identities remain private.


EXECUTIVE EXHIBIT 7 / ACTION PLAN

Make credit conditional, traceable and reversible

[PD] CMIC gives each open question an evidence artifact, an accountable owner and a decision gate. The proposed contract is designed to earn credit as well as to withhold it.

RELEASE CONTRACT / THREE GATES AND A REOPENING ROUTE

An open gate has an owner, an evidence requirement and a dated next decision



Accountable owner	Artifact and release condition
Metric and platform owners	A source-health and recovery map linking original cohorts, intake channels and classification versions. Reconstruct only where alternative evidence is validated; flag unresolved views.
Programme and evaluation owners	An authenticated treatment comparison with effect size, uncertainty, maturity and target-population bridge. Report credible positive findings as readily as adverse or null findings.
Finance and review owners	A reconciled claim register and a shadow-test decision log. Release credit only within the supported scope; record deferral cost and a reopening route.

[AI] The framework’s contribution is to separate three repair tasks: repairing the observation, identifying the intervention’s effect and governing the decision to recognize value. None substitutes for another. The falsifiable propositions specify where that separation should improve review decisions and where its cost may outweigh its value.

BOARDROOM DECISION

[PD] Set decision prices and an evidence budget before commissioning bounded shadow-review work. Preserve customer claim access while testing the programme’s benefits. Accept a benefit only at the population, treatment version and valuation basis that the evidence identifies.


[HR] Historical record in selected register S1; originals not recovered. [RC] Recalculation; [AI] analytical inference; [PD] proposed design. Historical quantities and identities remain private.

TECHNICAL ANALYSIS

Separate delivery, claim access, classification and credit

[AI/PD] The proposed Credit-and-Measurement Identification Contract (CMIC) prevents an administrative explanation from being mistaken for a complete causal verdict. A delivery policy, a claim-channel failure and a classification outage can operate simultaneously.

CMIC: CUSTOMER OUTCOME, ACCESS AND CLASSIFICATION ARE DISTINCT



U: baseline risk, route, geography and concurrent conditions affect A, Y and access; omitted arrows are explicit.

A randomized delivery-policy offer Z can identify its ITT only with stable, independently validated ascertainment.

The separate decision benchmark randomizes evidence workflow into E; it does not randomize historical delivery.

I is a selection variable. R can change at fixed Y. O is independently obtained, not independent of true Y.

[PD] A is a specified delivery-policy version. Y is physical nonreceipt at the original customer deadline, independently ascertained by O. C denotes attempted claims; I successful all-channel intake; F classification evidence; R the reported operational bucket. Channel conditions Q can change C, I and F without changing Y. The separate shadow benchmark randomizes evidence review, not A.

[AI] Baseline U includes address risk, route, customer access and concurrent conditions; it can affect policy targeting and Y. The policy can also change information, friction and claim behavior at fixed physical receipt. Conditioning on successful intake selects a population downstream of both customer need and access. A causal effect on R is not automatically an effect on Y.

[PD] Keep all original eligible deliveries, failed claim attempts and alternative channels in the reference frame. Fix treatment versions, clocks, geographic scope and risk definitions before analysis. Resolve reported classification mechanically where possible, while estimating customer and programme effects with a separate credible comparison. The arrows are proposed mechanisms, not historical findings.

Basis: proposed CMIC [AI/PD]; Sackett, Hernan-Hernandez-Diaz-Robins, VanderWeele-Hernan. Historical causal effects remain unidentified.

TECHNICAL ANALYSIS

A reported ruling does not settle every estimand

[HR/AI] The register records a challenge, escalation and a metric-owner explanation involving a failed classification feed. It does not supply an independently reproduced programme counterfactual or complete dated reconstruction.

Question	What would identify or close it?
Did the reported split change?	Stable event IDs, common population, immutable vintages and verified classification rules.
Did physical customer service change?	Independent receipt and original-deadline evidence on all original deliveries.
Did the programme cause a change?	Version-specific treatment assignment and a credible counterfactual with comparable ascertainment.
Can history be repaired?	A dated search of all surviving channel evidence, not absence in one source.
Can benefit be credited?	Causal effect, applicable volume, cost basis, overlap and uncertainty reconciled.

[HR/AI] The source reports unchanged total claims with a shifted operational split and names a platform mechanism. This can explain the classification movement. It does not prove the programme had no effect, eliminate a concurrent operating incident, or establish that blocked customers all reached another channel.

[HR/AI] A first decline near launch is not an invalid timing signature by itself. Similarly, missing original intake records do not establish permanent global irrecoverability: the later exchange leaves surviving signals under assessment. Distinguish service restoration, backfill feasibility and validated metric recovery.

[PD] Preserve the reported adjudication and unresolved questions side by side. Require independently evidenced onset, effective fix and data-comparability boundaries; retain unknown intervals. A calendar launch announcement does not establish actual exposure by region or version.

Basis: current registry and transcribed correspondence [HR]; identification critique [AI]. No historical quantity or private system identity is published.

TECHNICAL ANALYSIS

The same broken split admits opposite programme effects

[AI] This exact independent construction fixes 1,000 original deliveries. Mature recorded all-channel claims equal 100 before and after. Twenty lost exclusion records change the reported operational count from 60 to 80 in every world.

Synthetic post-period quantity	Beneficial	Null	Harmful
Observed claims with programme	100	100	100
Claims without programme (unobserved)	120	100	80
Programme effect: observed minus counterfactual	-20	0	+20
Recorded operational count before / after	60 / 80	60 / 80	60 / 80
Missing exclusion records	20	20	20

[AI] Let T1 be post-period claims under the programme and T0 the same population's unobserved post-period claims without it. The causal contrast is $(T1-T0)/N$, where N is original deliveries. The observed before-after contrast is zero; it does not determine T0. Exact binary potential-outcome assignments realize all three worlds, including the same actual feed break.

[AI] In a partial-information view, an analyst knows only that 20 of the 80 operational records lack exclusion evidence. The correct operational count can then span 60 to 80; the construction's hidden truth is 60. Narrowing that interval still does not identify T0. Even perfect repair of the reporting split cannot by itself establish programme credit.

[AI] These are claim outcomes, not independently verified physical nonreceipt. Claim attempts, successful intake and physical receipt must be separately measured. In a second construction, 100 physical failures stay fixed while successful claims fall from 80 to 40 solely through access loss; the apparent improvement is zero physical benefit.

[PD] The decision changes from choosing a single winning explanation to testing two contracts: repair or bound classification, and identify the programme's customer effect independently. Neither contract substitutes for the other.

Basis: exact synthetic potential-outcome and capture constructions [AI], counterexample.py/json. No value maps to historical coordinates.

TECHNICAL ANALYSIS

Shared baselines do not identify joint incremental benefit

[AI] Similar baselines are an additivity warning, not a mathematical proof of duplicate benefit. The missing quantity is the joint treatment outcome under a defined common population, version and financial basis.

Synthetic claim count	Additive effects	Redundant effects	Adverse interaction
Neither component Y00	100	100	100
Component A only Y10	80	80	80
Component B only Y01	80	80	80
Both components Y11	60	80	100
Interaction Y11-Y10-Y01+Y00	0	20	40

[AI] The two standalone reductions are identical across worlds, yet the joint reduction ranges from 40 to zero. The equation $\text{joint reduction} = \text{A-only reduction} + \text{B-only reduction} - \text{interaction}$ requires the joint cell. Population overlap alone does not fix the interaction, and similar starting rates do not prohibit genuinely incremental effects.

[HR/AI] The record mixes component entitlements, an annualized headline and a cumulative claim. Their arithmetic relation suggests a possible period explanation; it does not validate one. Nor does placing a worldwide claim beside a regional loss bucket make their ratio a commensurate coverage or duplication measure.

[PD] Reconcile each credit to version, counterfactual, geography, original volume, horizon, gross versus net basis and cost category. Use a bundle holdout or appropriately powered factorial contrasts for interactions. Component tests are not automatically additive; reference error and customer-access harm must enter the economic ledger.

[AI/PD] Released contact hours and route time are capacity, not cash until avoidability and realized use are demonstrated. A challenge's recorded contribution is separate from the programme benefit. Do not turn unverified claim crowding into a quantified amount of duplication.

Basis: independent factorial construction [AI]; source basis reconciliation [HR/AI]; Kohavi et al. on interaction and experiment trustworthiness.

TECHNICAL ANALYSIS

Validate the reference before interpreting a clean score

[PD] Specify separate measurement contracts for physical receipt, successful access to remedies, classification evidence and credit decisions. Administrative completion cannot be the sole reference for any of the others.

Construct	Required observation
Original exposure	Assignment, effective version, eligibility fixed before treatment, original volume and promise.
Physical receipt	Independent recipient/evidence adjudication of receipt by the original deadline; later recovery separate.
Claim access	Attempts, failed attempts, abandonments and cross-channel resolution on the original denominator.
Classification	Rule version, input presence, event time, posting time and each revision by immutable identity.
Decision quality	Prespecified credit rule, admissible packet, reviewer action, later blinded reference and effort.

[PD] Draw a prespecified probability sample across treatment, baseline risk, region, channel and incident exposure; include failed intake and no-claim deliveries. Two coders independently classify evidence while blinded to assignment and reported bucket. Report weighted confusion matrices, class prevalence, raw agreement and kappa, with domain-bootstrap intervals and independent disagreement adjudication.

[PD] A simple independent sample of 400 gives worst-case approximate 95% precision of plus or minus 4.9 percentage points for one binary agreement proportion, not for kappa or rare safety errors. Clustering, unequal weights and subgroup error rates require their own precision calculations. Test measurement invariance across channels and versions.

[AI/PD] Keep unknown references in original denominators. If at most epsilon in each original arm has unresolved binary truth, a contrast can move by at most 2*epsilon. This assumption sensitivity is not a missing-at-random claim. Timely capture is not correctness, and no record on one platform is not proof of no customer event.

Basis: proposed validation [PD]; Sackett, Cohen and selection literature. No historical classifier accuracy or independently validated receipt rate is inferred.

TECHNICAL ANALYSIS

Match populations and clocks without selecting away harm

[AI/PD] Matched weekdays, mature windows and incident segmentation improve definition but do not create a causal comparator. Excluding affected deliveries may change the target population or condition on consequences of treatment.

[HR/AI] The record contains programme before-after claims, targeted-stop cuts, an observational photo-timing comparison and a reported experiment in another region. Their denominators and treatment versions are not interchangeable. Reported statistical significance does not provide effect magnitude, customer value, durability or transportability.

[PD] Freeze the original cohort, assignment clock and outcome horizon. Analyze mature delivery cohorts rather than dividing newly posted claims by unrelated current volume. Estimate the full-cohort ITT first. A baseline-defined incident stratum can support prespecified heterogeneity; realized offline use, contact channel or delivery completion must not become an unnoticed eligibility filter.

[AI] Worst-performing targeted populations invite regression to the mean. Under an assumed linear conditional-mean model, untreated follow-up is $m + \rho \cdot (x - m)$, where x is the selected baseline, m the population mean and ρ assumed persistence. Hypothetical ρ values 0.4, 0.7 and 0.9 imply 60%, 30% and 10% reversion of selected excess. None is estimated here.

[RC/PD] The final inventory contains 19 source/sensitivity screens, 13 mechanical version audits, 5 design shapes and 10 profiles: 47 vectors. No fitted X/MR crossing occurs; none establishes a stable operating panel. For suitable ordered data, $MR_t = |x_t - x_{t-1}|$, X limits are $\text{mean}(x)$ plus or minus $2.66 \cdot \text{mean}(MR)$, and MR upper is $3.267 \cdot \text{mean}(MR)$. No eligible or computed Laney panel exists: transporter arm sizes are not binary outcome counts.

[PD] Prespecify baseline risk, device/network conditions, access needs, region, version and workload as moderators. Retain incident periods in sensitivity analyses. In a later observational ATT study, use compatible comparison cohorts, pretrend plots, equivalence and sensitivity assessments; nonsignificant pretrends are not proof of parallel trends.

Basis: source population/clock distinctions [HR/AI]; proposed analysis [PD]; Hernan et al., Cole-Stuart, VanderWeele-Hernan and Rosenbaum. Final source inventory is recorded privately.

TECHNICAL ANALYSIS

Credit the documented challenge within its evidence boundary

[HR/AI] The defensible contribution is a recorded challenge and escalation about attribution. Programme design, experiments, rollout, official root-cause adjudication and downstream remediation have different owners.

Claim	Evidence test and bounded interpretation
Questioned attribution and escalated	Provisional registry-level hoop: transcribed outgoing messages support a challenge; originals and timestamps require authentication.
Specified a complete causal standard	The richer mechanism and matched/mature/segmented rule appear in derivative working notes; do not present them all as verbatim sent text.
Caused the substantive reply	Sequence is compatible with influence; the reply followed escalation, but no counterfactual isolates the escalation's effect.
Solved the cause or earned programme savings	Unsupported here. The metric owner supplied a reported ruling; this record does not establish the challenger's programme design or evaluation.
Obtained operational closure	Incomplete: fix and comparability dates and surviving-signal reconstruction remain unverified.

[HR/AI] Absence from a contributor table bounds what this record supports; it does not establish universal absence of participation. A copied recipient is not an analytical contributor. Separate source authentication, interpretation and our present independent counterexamples in the contribution ledger.

[PD] Apply process tracing to alternatives: routine investigation already under way, parallel escalation, reviewer incentives and delayed response. Seek signed requests, intermediate analyses, delivery receipts and acceptance records. Retrospective narration and repeated copies from the same thread are not independent corroboration.

[PD] Close the dispute with a dated evidence owner for service restoration, historical reconstruction, metric certification and benefit validation. A reporting caveat can be warranted before those tasks close; it must not be promoted into a certified causal-null result.

Basis: current registry-mediated role record [HR]; contribution and process-tracing critique [AI/PD]. No sole authorship, deployment or realized benefit is claimed.

TECHNICAL ANALYSIS

Design a prospective test that can reject a favorable story

[PD] The evaluation must separately test the decision process and any delivery programme. A benchmark of credit adjudication can establish better decisions under its evidence law; it cannot establish physical customer benefit from the historical rollout.

[PD] Freeze all original dockets, resource prices and reference rules before domain assignment. Lock decisions at intake + 48 hours and complete blinded adjudication by intake + 14 days. Preserve equal evidence rights and current customer treatment. Shared reviewers or source faults belong in one domain; dockets are conditionally independent only given the modeled shared reference condition.

[PD] Let $L(z)$ be prespecified false-credit, missed-benefit and defer/rework loss under offer z , plus variable verification work. Its ITT is $E[L(0)-L(1)]$ per original docket. Let K be fixed incremental production cost; funding requires a lower bound above K and a separate lower bound above zero for usual-minus-CMIC false credit. Known mandatory failures defer; unknown reference cases remain in bounds.

[PD] The numerical physical-contrast signals are constructed evidence under known synthetic programme states. They do not imply that an individual receipt observation reveals a causal effect. Reviewers see only permitted signals; the hidden final reference must not leak into policy tuning or adjudication at the decision deadline.

[PD] For a distinct delivery-policy trial, use independently reconstructed original-deadline receipt and legitimate claim access on all original deliveries, alongside full operating resources. A lower recorded claim count or a repaired operational split cannot substitute for either endpoint. Region/version extrapolation requires a separate transport analysis.

[PD] The numerical design on the following pages is wholly hypothetical. Independently reproduce its finite law, exact moments, boundary-specific covariance, economic break-even and directional detection. Compare MDE with the economically relevant effect, test the complete joint decision rule, and simulate actual boundary laws with Monte Carlo intervals. Approximate normal planning is not an exact finite or all-nuisance guarantee.

[PD] Prespecify stopping, multiplicity, adjudication and negative-control probes. A probe must share the relevant measurement failure while lacking the programme mechanism; its result neither quantifies bias nor certifies every pathway. Recruitment, reference quality and protocol feasibility must be calibrated before use.

Basis: prospective CMIC [PD]; Kohavi et al., Lipsitch et al. and independent planning. Numerical benchmark is a proposed design, never historical impact.

TECHNICAL ANALYSIS

Make the identification contract falsifiable

[PD] CMIC is proposed. Each proposition has an observable failure criterion; successful documentation alone cannot confirm a customer effect or monetary credit.

Proposition	Test and boundary
P1: Separate pipeline checks improve diagnostic decisions	With independent truth and fixed cases, classification checks plus a separate customer-effect test should reduce false causal closure. Failure to reduce the prespecified error rejects this protocol benefit.
P2: Access-complete reference improves decisions when capture differs	Prespecified multi-channel audits must reduce classification/credit error against blinded reference, not merely increase records. Failure to improve that endpoint rejects the benefit claim.
P3: Version-compatible transport predicts target effects	A target-population replication must match the prespecified transported effect interval. A version change or failed overlap check prevents that test from certifying the old claim.
P4: Credit validation earns its cost without unsafe decisions	The complete offered protocol must clear economic and harmful-credit gates. High agreement or fewer disputes with no net decision value rejects the funding case.

[AI] Deming warns against reacting to a measurement break as operating variation. Campbell/Goodhart predicts adaptation when claim counts become targets. Ittner-Larcker requires testing the chain to value; Tucker-Edmondson focuses attention on workarounds such as alternate channels. Composite-measure principles require explicit weights, omitted harms and nonoverlapping cost categories. CMIC extends these ideas by keeping customer truth, access, classification and credit as separately testable contracts.

Limitations and identifying assumptions

[AI/PD] Original experiment and correspondence artifacts, linked mature customer outcomes and authenticated effective versions remain unavailable in this register. The exact examples prove nonidentification, not a historical mechanism. Proposed benchmarks require calibrated evidence accuracy, prices, deadline loss and independent domain sampling. Field inference additionally needs stable reference quality, treatment consistency, overlap and controlled interference. A reporting repair can succeed while programme credit remains unresolved; prospective validation does not retroactively identify it.

Basis: proposed propositions [PD], independent intellectual contribution [AI] and explicitly bounded literature applications. No historical causal effect is certified.

Credit requires two independent questions

[PD] Test CMIC as a shadow review workflow. It must distinguish physical service benefit from impaired access to legitimate claims. This benchmark changes no delivery policy, customer entitlement or accounting credit.

[PD] Randomize independent reviewer/source domains to usual headline review or CMIC. Each receives 24 frozen dockets, independent conditional on shared domain grade, from the same eligibility law, with equal source access and authority. Lock credit, reject or defer by original intake + 48 hours; blind adjudication by intake + 14 days retains every docket and is hidden from rule tuning and 48-hour decisions. Shared teams or source faults belong in one domain.

Synthetic primitive	Locked interpretation
Hidden states: 45% / 35% / 20%	Sound benefit / no physical benefit / benefit with impaired legitimate claim access
Usual: six reported contrasts	Credit a positive classification sum; otherwise reject
CMIC: 12 physical contrasts + eight access probes	Credit physical sum >= 2 and access failures <= 1; reject sum <= -2 or failures >= 3; otherwise defer
Known mandatory blocker: 10%	Defer in both arms; never infer a missing check passed
Reference corruption / access sensitivity	Sound grade: 5% / 95%; fragile: 60% / 25%
Decision loss: credit / reject / defer	Sound state: 0 / 12 / 6; other states: 12 / 0 / 6 resource units

[PD] Signal likelihoods are invented and fallible. A shared fragile-reference condition affects all dockets. Physical contrasts are constructed experimental evidence, not individual causal observations; linkage and access checks remain fallible. Hidden states are unavailable to reviewers. The source supplies none of these probabilities.

[PD] THE TARGET IS DECISION QUALITY

False credit means approving a programme without both sound physical benefit and preserved legitimate access. Correct classification, sponsor agreement or fewer recorded claims cannot substitute for these conditions.

[PD] Constructed shadow cases, not historical observations or customer-effect estimates. Reference correctness must be calibrated.

PROPOSED CREDIT-ASSURANCE EVALUATION / 2 OF 4

Power the value and false-credit gates

[PD] Primary gain is usual minus CMIC mean decision loss and verification work. The second gate requires fewer false credits per all assigned dockets. Positive benefit detection must also reach 80% at economic break-even.

[PD] Exact finite signal moments include the persistent domain condition. Normal inference uses a variance envelope over a prespecified response mixture: lambda invokes CMIC rather than usual review per docket. Lambda is not observed compliance. Fixed capacity is paid even at lambda = 0.

Direct design	Domains / arm	MDE (index)	BE detection	Joint floor
Small check	16	> full	22.1%	0.0%
Base candidate	384	4.71	>99.9%	91.3%
Fragile stress	384	5.16	>99.9%	9.2%
Corroborated version	128	8.72	>99.9%	>99.9%

[PD] Base fixed cost is 10.11 index points. The stress raises fragile-reference prevalence from 15% to 60%. The redesign improves corroboration and access-failure detection; its cost is 15.62 points. It is a new workflow version, not a sample-size cure.

[PD] Economic clearance requires the lower 97.5% one-sided normal bound to exceed fixed cost; false-credit superiority requires its own lower bound above zero. MDE uses the beneficial direction of a two-sided 5% zero test; high power against zero does not imply economic clearance. The joint normal-planning floor sums marginal powers minus one; zero is uninformative.

[PD] CONDITIONAL RECRUITMENT ONLY

The base allocation entails 768 independent domains and 18,432 dockets. No historical catalogue establishes this recruitment or the assumed likelihoods. Start with bounded reference/variance calibration; recalibrate before operational use.

[PD] Usual base modeled resource burden = 100; MDE and costs are index points. Normal coverage and power are approximate, not finite guarantees.

Price useful decisions, not suppressed claims

[PD] All monetary-looking weights below are hypothetical resource units. No refund reduction, programme entitlement, recovered revenue or historical savings enters the calculation.

[PD] Verification work is 0.25 per invoked docket, including known blockers and wrong decisions. Fixed cost is 0.55 per assigned base docket: 0.20 review capacity, 0.20 reference governance and 0.15 coordination. Corroboration adds 0.30. These production costs exclude research-only blind adjudication, which needs a separate study budget.

[PD] For a declared state, reference stratum and loss-category comparison, let F be favorable expected decision-loss change and H adverse change. $Net_raw = a * h * F - H - \text{verification work} - K$. Here a is avoidability and h is non-overlap; added burdens remain fully priced. This conditional-mean decomposition is a hypothetical valuation convention, not identified individual causal savings.

Scenario / realization	a	h	Net index
Base / full	1	1	5.48
Base / partial	0.8	0.9	-10.15
Fragile / full	1	1	1.16
Redesign / full	1	1	15.70
Redesign / partial	0.8	0.9	-5.11

[PD] Prices must be set independently before validation. At defer price 3, always-defer costs 3 while full CMIC costs 3.87; it dominates. At price 6, usual / CMIC / always-credit / always-reject / always-defer cost 5.44 / 5.14 / 6.54 / 5.46 / 6.00. At defer price 7.5, usual review wins. Price 6 is uncalibrated; changed prices require re-powering, not reuse of the base power.

[PD] Resolve the original decision state using all channels and the original approved entitlement rule. Later receipt does not rewrite the earlier state. With up to epsilon unresolved or misadjudicated dockets in each arm, loss contrast can shift by $24 * \text{epsilon}$ and false-credit contrast by $2 * \text{epsilon}$; work remains separately recorded. Reference errors are not missing-at-random by assumption.

[PD] CUSTOMER VALUE REQUIRES A DIFFERENT TEST

A later randomized delivery-policy trial would estimate original-cohort physical and access ITTs under equal remedies. Its valuation must price avoided physical harm minus added access harm and full work. This shadow benchmark provides no such effect or cash estimate.

[PD] $Net_index = 100 * Net_raw / \text{base usual burden}$. Scenario contrasts and resource prices are invented; endpoint reference bounds use all original dockets.

PROPOSED CREDIT-ASSURANCE EVALUATION / 4 OF 4

Buy evidence only when it can change the decision

[PD] A separate one-shot information model asks whether to authorize bounded assurance, repair the reference first or defer investment. Its utility is not added to the benchmark gain.

Action / latent state	Decision-ready	Repairable	Unresolved
Authorize assurance	12	-8	-16
Repair reference first	4	7	-2
Defer investment	0	0	0

[PD] Hypothetical state probabilities are 40% / 35% / 25%. Current best expected utility is 3.55; perfect information can add at most 3.70. Signal accuracy is q , with each wrong label $(1 - q) / 2$. Evaluate all 27 policies before subtracting cost.

Evidence quality	q	EVSI	Cost	Net EVSI
low	0.45	0.0000	0.40	-0.4000
base	0.70	0.9275	1.20	-0.2725
high	0.90	2.6900	2.80	-0.1100

[PD] Six fresh finite-process checks use 12,000 trials each. Full-response and economic-boundary results below have Wilson 95% intervals. Zero/false-credit boundaries are also checked privately; these selected laws do not establish universal coverage.

Finite simulation	Joint success	False economic clearance
Base	91.22%	2.23%
Corroborated	99.96%	1.74%

[PD] Joint intervals: base 90.70-91.71%, corroborated 99.90-99.98%. Economic-boundary intervals: 1.98-2.50% and 1.52-1.99%. These measure Monte Carlo uncertainty, not field parameter uncertainty.

[PD] RELEASE GATE

Fix scoring, prices and thresholds before validation. A feed fault can coexist with a real programme effect. Deferring credit cannot justify withholding a customer remedy.

[PD] Separate decision-utility units. $EVSI \leq EVPI$; evidence purchase does not authorize deployment or establish historical benefit.

REPRODUCIBILITY / READER PROTOCOL

Reproduce the distinction before releasing the credit

[AI] CMIC separates the customer outcome, its observation, its causal explanation and the decision to recognize programme value. Each layer needs a different evidence artifact.

Reproduction layer	Required reconstruction
Source and contribution	Compare selected cells with the immutable register. Preserve original-recovery failures, quoted messages, narrative reconstructions, disputed clocks and the division of decision authority.
Measurement and identification	Rebuild the classification and claim-access constructions. Verify original denominators and the assumptions supporting the counterfactual comparison. Changing the reporting split must not become a causal-null claim.
Scope and valuation	Reconcile component precision, reporting horizons, geography and treatment versions. Reproduce the factorial worlds before assuming component effects are additive or duplicated.
Evaluation and economics	Rebuild the finite review law, fallible evidence, frozen decision policies, domain dependence and complete resource loss. Compare against simple review alternatives and fully charge deferral.
Decisions and disclosure	Recompute the direct allocations, selected-boundary simulations, valuation sensitivity and separate evidence-acquisition decision. Match published conclusions and reviewer records to the final file hashes.

Protected history and independent constructions

[AI] Historical numerical coordinates, identities and version conflicts remain in the separate private annex. Public numerical examples are explicitly hypothetical and source-independent. A common positive rescaling preserves effects and decision signs only when all corresponding costs and thresholds use the same scale.

Meaning of a future result

[AI] The proposed shadow trial estimates the effect of a specified review package on decision-resource loss and false credit in eligible independent review domains. It does not estimate the physical effect of the historical delivery programme. Blinded reference evidence must be validated; latent truth in a simulation is not automatically observable in practice.

PRACTICAL HANDOFF

[PD] Validate the measurement path, authenticate the comparison and reconcile the value basis. Set a loss schedule and evidence budget before commissioning the shadow test. Reopen credit decisions when credible new evidence changes the supported contrast.

Private source maps, exact calculations, finite model and hash-bound reviews support reproduction. Reproducing a hypothetical model does not authenticate historical programme impact.

RESEARCH REFERENCES / CAUSAL IDENTIFICATION

Separate the outcome, its observation and its cause

Primary research informs interpretation and prospective design. It does not authenticate historical artifacts, establish adoption of a method or supply the project's missing causal effect.

Hernan, M. A., Hernandez-Diaz, S., & Robins, J. M. (2004). A structural approach to selection bias. *Epidemiology*, 15(5), 615-625. doi:10.1097/01.ede.0000135174.63482.43. Primary paper read. Application: selection on a downstream variable can open a noncausal path; original cohorts and downstream reporting cuts are distinct.

Sackett, D. L. (1979). Bias in analytic research. *Journal of Chronic Diseases*, 32(1-2), 51-63. doi:10.1016/0021-9681(79)90012-2. Original article copy read. Application: detection and response mechanisms can change apparent outcomes; validate physical receipt and claim access separately.

Cole, S. R., & Stuart, E. A. (2010). Generalizing evidence from randomized clinical trials to target populations: The ACTG 320 trial. *American Journal of Epidemiology*, 172(1), 107-115. doi:10.1093/aje/kwq084. Primary full text read. Application: trial validity does not establish target-population transportability without the relevant population bridge.

VanderWeele, T. J., & Hernan, M. A. (2013). Causal inference under multiple versions of treatment. *Journal of Causal Inference*, 1(1), 1-20. doi:10.1515/jci-2012-0002. Indexed primary-text sections accessed; full PDF review is not claimed. Application: a shared programme name does not establish identical intervention versions.

Lipsitch, M., Tchetgen Tchetgen, E., & Cohen, T. (2010). Negative controls: A tool for detecting confounding and bias in observational studies. *Epidemiology*, 21(3), 383-388. doi:10.1097/EDE.0b013e3181d61eeb. Primary full text read. Application: prespecified negative controls can reveal particular biases; passing a check does not prove all confounding absent.

Kohavi, R., Deng, A., Frasca, B., Walker, T., Xu, Y., & Pohlmann, N. (2013). Online controlled experiments at large scale. *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1168-1176. doi:10.1145/2487575.2488217. Original conference paper read. Application: instrumentation, interactions and experimental integrity require explicit checks before crediting changes.

Rosenbaum, P. R. (1984). The Consequences of Adjustment for a Concomitant Variable that Has Been Affected by the Treatment. *Journal of the Royal Statistical Society: Series A (General)*, 147(5), 656-666. doi:10.2307/2981697. Publisher abstract and metadata accessed. Application: adjusting for treatment-affected offline behaviour or contact channel can bias the comparison; no historical bias magnitude is inferred.

Verified primary scholarly or official sources. Access scope is disclosed in the private research record; cross-domain methods do not transfer empirical effects.

RESEARCH REFERENCES / EVALUATION AND MEASUREMENT

Test the review at its independent decision level

Primary research informs interpretation and prospective design. It does not authenticate historical artifacts, establish adoption of a method or supply the project's missing causal effect.

NIST/SEMATECH. e-Handbook of Statistical Methods, section 6.3.2.2: Individuals Control Charts. Official method. Calculate tentative limits for the sparse ordered observations, but separate planned milestones, cross-sectional groups and conflicting source definitions. Within-limit movement does not establish either improvement or a stable operating phase.

Laney, D. B. (2002). Improved Control Charts for Attributes. *Quality Engineering*, 14(4), 531–537. Primary publisher abstract. Extra-binomial adjustment requires actual counts, denominators and rational subgroups. The available baseline snapshots lack an eligible repeated count panel; rounded rates are not converted into invented exposure counts.

Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20, 37–46. Primary publisher record. Assess agreement and kappa in a blinded, stratified taxonomy audit. Report prevalence, confusion and uncertainty; high agreement between coders cannot authenticate a shared defective reference.

Barnett, A. G., van der Pols, J. C., & Dobson, A. J. (2005). Regression to the Mean: What It Is and How to Deal with It. *International Journal of Epidemiology*, 34, 215–220. Primary abstract. A selected worst-site cohort can recover without an intervention. The case uses explicitly assumed reliability sensitivity and requires independent recurrence and comparison data.

Hudgens, M. G., & Halloran, M. E. (2008). Toward Causal Inference with Interference. *Journal of the American Statistical Association*, 103, 832–842. Original manuscript. Analysts and decision artifacts can transmit a review policy across assigned units. Group shared reviewers into independent review domains and assess cross-domain leakage before treating assignment as isolated.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, 225, 200–230. Author manuscript. A later nonrandom phased review rollout needs group-time effects, eligible comparisons, parallel-trend diagnostics and its own power analysis. It does not retroactively identify the historical programme effect.

Strong, M., Oakley, J. E., Brennan, A., & Breeze, P. (2015). Estimating the Expected Value of Sample Information Using the Probabilistic Sensitivity Analysis Sample. *Medical Decision Making*, 35, 570–583. Author repository. Information has value when it changes a feasible decision under a specified prior, payoff and likelihood model. Evidence cost is subtracted; EVSI is not added to the trial's operating benefit.

Verified primary scholarly or official sources. Access scope is disclosed in the private research record; cross-domain methods do not transfer empirical effects.

RESEARCH REFERENCES / GOVERNANCE AND INFERENCE

Connect accountability to a validated causal claim

Primary research informs interpretation and prospective design. It does not authenticate historical artifacts, establish adoption of a method or supply the project's missing causal effect.

Collier, D. (2011). Understanding Process Tracing. *PS: Political Science & Politics*, 44, 823–830. Primary publication. Grade each specific role and action against competing explanations. A derivative register supports a provisional contribution claim; a named request, working draft or colleague's delivery row does not establish convening, delivery or sole authorship.

Tucker, A. L., & Edmondson, A. C. (2003). Why Hospitals Don't Learn from Failures. *California Management Review*, 45(2), 55-72. Author manuscript. Local workarounds can resolve immediate needs while hiding system conditions. CMIC links source repair, causal evidence and financial recognition to accountable owners; it does not transfer a healthcare effect.

The W. Edwards Deming Institute. The Funnel Experiment. Official exposition. Repeated adjustment to routine variation can worsen performance. The control-assurance mechanism therefore requires stable measurement, recurrence evidence and a prospective comparison; sparse dashboard fluctuations do not justify repeated changes to frontline instructions.

Campbell, D. T. (1976). Assessing the Impact of Planned Social Change. Original institutional record. Accountability measures can change both behavior and reporting. Track eligibility, capture, unknowns and independent service outcomes so that changing an exemption or reason code cannot substitute for improving physical receipt.

Goodhart, C. A. E. (1992). The Objectives for, and Conduct of, Monetary Policy in the 1990s. Reserve Bank of Australia Annual Conference. Original author paper, section 1(b)(i). Relations between intermediate targets and ultimate objectives can change with adaptation. Apply that mechanism to recorded claim labels and physical customer outcomes; no monetary-policy effect is transferred.

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. Strategic measures require an explicit mechanism and validated connection to outcomes. A corrected classification does not establish the causal or economic value of a programme.

OECD / European Commission JRC (2008). Handbook on Constructing Composite Indicators. Official handbook. Definitions, denominators, weights and compensation assumptions govern aggregation. Keep customer access, claim classification, decision value and false-credit safety distinct; a favorable composite cannot compensate for a failed identification requirement.

Verified primary scholarly or official sources. Access scope is disclosed in the private research record; cross-domain methods do not transfer empirical effects.

Accept the identification correction and the bounded review test

[AI] This author-led, agent-assisted review is not external peer review or client approval. Acceptance applies to evidence reconciliation and proposed evaluation, not verified programme impact.

Major issue	Disposition in the revised case
Diagnosis mistaken for causal null	Three exact worlds retain the same observed claim total and broken split while programme effects differ in sign. The recorded ruling remains a recorded interpretation.
Customer access hidden by the metric	Separate physical failure, legitimate claim completion, capture and classification. A fall in reported claims cannot by itself establish improved service.
Financial discrepancy overinterpreted	Retain rounding-compatible totals while requiring a valuation-horizon bridge. Geography, component interactions and paid-capacity conversion remain separate identification tasks.
Personal impact overstated	Retain provisional authorship of the challenge and follow-up. Do not claim programme delivery, causal acceleration, prevented booking or permanent historical irrecoverability.
Review model treated as field evidence	Use a fallible finite evidence law, explicit resource prices, competing simple policies, direct sample allocations and selected-boundary checks. Require independent reference and loss-schedule validation.

Partner critique: can a COO act on page one?

[AI] Yes: protect the affected metric from unsupported causal credit, assign the missing evidence to its owner and test the review package before scaling it. A reporting fix and a financial sign-off are separate decisions. The case supports disciplined challenge without claiming ownership of another team’s programme.

Remaining conditions and minor issues

[AI] Original authentication, cross-channel reconstruction, treatment matching, loss-price calibration, independent-domain recruitment and reference validity remain open. The hypothetical model is conditional planning evidence. No historical savings, external endorsement or consulting-firm ranking is established.

REVIEW DISPOSITION / PASS WITHIN STATED SCOPE

[AI] Accept the reconciled evidence case and falsifiable review design. Failed measurement, an unidentified contrast, inadequate economic detectability or failed false-credit protection blocks expansion under the proposed decision rule.

Internal source, methodological, numerical, executive and actual-page review. No external referee endorsement or verified client effectiveness is claimed.