

Program Portfolio & Ownership Governance

A reconciled row is not an approved programme

An operating-project case on rebuilding an annual-plan intake register. It separates source-field provenance, proposal status, accountable ownership and joint benefit evidence from the authority to fund and the ability to deliver.

PORTFOLIO INTAKE & DECISION INTEGRITY

BLUF / BOTTOM LINE UP FRONT

Make every row traceable. Keep funding and joint value as separate decisions.

ANALYTICAL FINDING

Common populations and identical estimates flag a question; they do not measure duplicated benefit. A complete register can preserve uncertainty while a committee still lacks the evidence to fund a bundle.

CONTRIBUTION IN THE SOURCE RECORD

Provisional registry-level hoop: reported rowset construction and reconciliation; no ownership of funding or programme results.

EVIDENCE STATUS

Evidence-reconciliation and proposed-design case. Original artifacts remain unauthenticated; downstream closure and impact are unestablished.

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

STRUCTURED ABSTRACT

Prepare an auditable decision packet without declaring the decision complete

[AI] The register improves visibility when it states what is known, what was read and what remains owed. Those states must survive the move into a funding discussion.

Element	Decision-relevant conclusion
Context [HR]	A planning lead reportedly rebuilt a filed intake view alongside additional internal proposals, aligning source assignees, value types, confidence labels, metrics and outstanding validation.
Objective [AI]	Preserve the distinction between an accurate transcription, an accepted accountable role, a filed request and an authorized, executable programme.
Approach [RC/AI]	Reconcile row populations and field definitions; retain missing-read states; inspect benefit relationships without manufacturing a net total from aggregate summaries.
Findings [AI]	A changed owner field does not prove the prior role was wrong. A generic metric picker can coexist with a named follow-up answer. A common measured pool does not identify joint benefit.
Decision [PD]	The intake owner obtains a current source snapshot and accepted role definitions. Programme owners and Finance reconcile joint value and scope; the committee records the actual funding ruling.
Implications [PD]	Use the proposed Row-to-Decision Evidence Contract (RDEC). Test a shadow packet policy against fixed choices and fuller review before buying it; synthetic precision alone cannot justify deployment.

[PD] A packet may be complete enough to expose its open questions while still being ineligible for funding. Require a valid filed identity, accepted decision authority, dependency evidence and an explicit treatment of joint value. Preserve unsized and unread states without converting them into zero economic value.

[HR] historical record; [RC] recalculated; [AI] analytical inference; [PD] proposed design. Protected source S1. Historical identities, quantities, indexed coordinates and internal locators remain private.

EXHIBIT 1 / ROW POPULATION

A fuller intent register must not silently become a larger filed ask

[AI] The internal programme list and the submitted intake are different sets. Combining them helps planning only when their eligibility and decision states remain separate.

Row state	What the record establishes	What is still required
Filed request	A source-linked intake record exists at a stated vintage.	Current committee ruling, conditions and dependency acceptance.
Internal proposal	An intended programme is visible in the planning view.	A filed entry or explicit exception under the governing process.
Closed or duplicate status	The source carries a status with that wording.	Authority, effective version and whether the closure applies to this exact request.
Confidence assessment	An editor has rated delivery confidence.	Independent evidence for funding, capacity and the underlying assumptions.

[HR/AI] The historical rowset intentionally included proposals absent from the filed board. That resolves the apparent row-count disagreement as a population difference; it does not authenticate a current complete dataset or show that the additional requests were submitted. A general board link also does not substitute for an auditable row-level identifier.

[PD] Publish separate views for intent, filed requests and approved scope. Reconcile additions, removals and changes by stable identity and effective date. Make incompatible scopes impossible to combine through an unqualified total, rather than relying on a reader to notice a caveat in a long notes column.

KEEP THE GATE VISIBLE

[PD] The register owner controls representation. The committee controls approval. An internal confidence label cannot silently replace either the source status or an accepted funding decision.

[HR/AI] Source aggregate descriptions are not the complete original row-level dataset. No public historical totals, distributions or invented cross-tabulations are used.

EXHIBIT 2 / FIELD MEANING

Type the value and the read state before reconciling the number

[AI] Numeric-looking cells and blank fields can encode different constructs. Arithmetic can close while their meanings remain incompatible.

Evidence state	Safe representation	Unsafe inference
Sized owner estimate	Amount, unit, scope, vintage and validation owed.	Finance has approved it or it will be realized.
Unsize marker	Explicit unsize category under a confirmed local convention.	Literal money, a zero estimate or proof of no benefit.
Explicit zero	Zero for the named measure and scope, with reason.	Zero everywhere, including another business scope.
Field not captured	Unread, with the missing source panel identified.	The source itself has no answer or the answer is No.
Read but absent answer	Missing at the checked source and vintage.	The programme or dependency does not exist anywhere.

[HR/AI] The record describes a local sentinel convention alongside separately represented zeroes and uncoded values. That supports a contextual reading rule, not a universal proof from the coexistence of two numbers. Confirm ambiguous use with the submitting owner. A cost field can also contain a placeholder; the literal sum is not automatically a funded resource requirement.

[HR/AI] A headline-total difference reconciled to one cross-business entry is an arithmetic bridge. Unread component fields prevent that equality from proving complete coverage or accepted value.

TURN NOTES INTO OPERATING FIELDS

[PD] Keep source value, interpreted type, read coverage, scope and validation state separately. Preserve the raw entry for audit. A warning beside a numeric cell is insufficient if downstream software can still sum it as an eligible amount.


[AI/PD] Proposed typed representation. An unresolved Finance flag or uncaptured panel prevents a verified-approval claim; it does not establish that approval or value is zero in every source.

EXHIBIT 3 / EVIDENCE LAYERS

Trace the row through the decision without skipping the missing links

[AI] Source capture, a reconciled record, a committee ruling and mature outcomes are distinct evidence layers. Passing one does not certify the next.

EACH TRANSITION NEEDS ITS OWN EVIDENCE



NO AUTOMATIC PROMOTION
A filled cell cannot approve funding or establish an operating effect.

[HR/AI] The record reports parser shape checks, identity resolution, audits of the written workbook and a claim verifier. These can catch transcription or build defects. Their described behaviour is not a replay of the original scripts, and a correct derivative does not establish that the live filed board was repaired.

[PD] Reconcile both the written output and the current source of record. Retain row identifiers, version hashes, transformation rules, exceptions and reviewer acknowledgements. Then capture the committee decision and the conditions that must hold before execution.

A GAP REGISTER IS AN OUTPUT, NOT CLOSURE
[PD] Assign each unresolved item to an accepted owner with evidence, due date and reopening trigger. Close it only against the relevant source or decision record. A complete punch-list can coexist with every substantive dependency still open.

[PD] Conceptual diagram without historical quantities. Provenance supports tracing the claim; it does not itself validate economic truth, decision authority or causal delivery benefit.

EXHIBIT 4 / OWNERSHIP AND MEASUREMENT

Align role definitions before calling a changed owner a correction

[AI] A source assignee, an internal coordinating lead and an accountable delivery owner can be different people. Matching the chosen source field is a provenance result, not an acceptance test.

Question	Required evidence	Decision boundary
Who is recorded?	Source field, identity resolution and checked output.	Confirms the recorded assignee at that vintage.
Who is accountable?	Accepted scope, role definition and delegation.	A submitter address alone remains provisional.
Who can deliver?	Capacity, dependencies and authority for the work.	A name or confidence rating is insufficient.
What is measured?	Named construct, formula, unit, cohort and source.	A picker value and its free-text follow-up must be read together.

[HR/AI] The owner convention changed from an internally inferred lane role to the assignee recorded in the source. Many changed cells therefore show definition alignment; they do not prove that every prior role assignment was factually wrong. Missing assignees filled from submission identity still require confirmation.

[HR/AI] Generic metric choices already had named follow-up answers in the reviewed descriptions. The rewrite improved their visibility and stated baseline validation owed; it did not demonstrate that those programmes previously had no metric. Likewise, a metric identifier is a locator, not proof of a valid denominator or an approved target.

BOUND THE CONTRIBUTION

[HR/AI] Credit the planning lead with reported artifact construction, editorial rules and reconciliation, provisionally at registry-level hoop strength. Do not transfer programme ownership, Finance validation or committee decisions to the person compiling the register.

[HR/AI] Role and metric semantics are taken from the reviewed descriptions, with contradictory interpretations retained privately. No inference about individual performance is made.

EXHIBIT 5 / JOINT BENEFIT

A shared population can imply duplication, addition, complementarity or harm

[AI] The same standalone improvements do not identify what happens when both programmes operate. A common population is a reason to investigate their joint effect.

Independent constructed world	Baseline	A alone / B alone	Both	Joint benefit
Redundant	100	70 / 70	70	30
Additive	100	70 / 70	40	60
Complementary	100	70 / 70	10	90
Interfering	100	70 / 70	100	0

[PD] Lower outcomes are better in this fixture. Every world has the same population and standalone benefit of 30 for each programme. Yet the combined benefit ranges from zero to 90. The example fixes potential outcomes by construction; it is not an estimated interaction or a statement about the historical programmes.

[HR/AI] Shared measured pools, identical amounts and similar physical control points were recorded as overlap exposures. Those signals do not quantify overlap, establish that one sizing was copied, or authorize merging distinct requests. Opposite differences from separate planning estimates do not resolve that question either.

ASK FOR THE JOINT CONTRAST

[PD] Define programme versions, eligible units, shared resources and the no-programme baseline. Estimate or bound the combined effect and full combined cost. Do not add standalone effects automatically; do not subtract one merely because the denominators match.

[PD] Source-independent synthetic potential-outcome construction, independently reproduced. The source has no linked joint-treatment outcome panel sufficient to identify these effects.

EXHIBIT 6 / PORTFOLIO READINESS

Expose missing costs and accepted dependencies without manufacturing a causal story

[AI] A benefit register can fit its funded headcount while omitting a cost class or an external commitment. Arithmetic reconciliation is weaker than executable delivery.

Observed register feature	Question to resolve	Claim to avoid
Headcount fits a funding envelope	What work, recurrence and calendar capacity were included?	The allocation proves productivity or full feasibility.
Fixed costs have no destination column	Where are capital, setup and non-labour costs authorized?	Benefit entitlement funds those costs automatically.
External dependency says Pending	Which version, team and delivery window are accepted?	Text in a link field is a committed dependency.
Unpriced safety-related proposals carry unfavorable output ratings	What were the actual decision criteria and feasible alternatives?	Lack of a benefit estimate caused the decision.

[HR/AI] The reviewed account contains confidence assessments, technology positions and committee-status labels that belong to different decisions. It also describes value flags that conflict with explanatory notes. Preserve the conflict and obtain the relevant authority record; neither a headline Yes nor a conservative recognition rule settles all underlying evidence.

NO UNOBSERVED CORRELATIONS

[PD] Recover the authoritative rowset before analysing how confidence, funding, cost or ownership relate. The selected register descriptions do not support fabricated row correlations, programme rankings or a causal explanation for which requests were deferred.

[HR/AI] Historical aggregates and exceptions remain in the private annex. The present case identifies missing decision evidence, not realized savings, funded capacity or the cause of committee choices.

EXHIBIT 7 / DECISION CONTRACT

Test the reconciliation policy against the choices already available

[PD] The proposed Row-to-Decision Evidence Contract keeps valid source information, substantive joint value and authorization separate while making the next decision explicit.

RDEC obligation	Release evidence
Preserve population and status	Original proposal identities; filed/internal scope; known authority blockers remain visible.
Type and trace each field	Read state, meaning, source, version and accepted owner convention.
Test joint value	Standalone and combined mechanisms, resource interaction, costs and uncertainty.
Retain strong alternatives	Best fixed recommendation, safe fallback, unsized response and fuller review on matching evidence rights.
Close the actual decision	Accepted committee ruling, dependency commitments and later independent outcome evaluation.

[PD] The synthetic shadow benchmark tests recommendations without allocating real funds. A safe single-package fallback avoids unnecessary deferral when joint evidence is unresolved. The small and ambiguous-evidence studies lack break-even precision. Better evidence and a larger candidate improve conditional power, but fuller review has lower total loss at the declared prices.

NEXT ACTION


[PD] Recover a bounded current-source sample, confirm role and sentinel conventions, and resolve one joint-benefit question with the relevant owners. Measure review effort and decision errors before a larger replay. Keep the original proposals in view; do not call record cleanup a funded portfolio or a delivered customer benefit.

[PD] The benchmark and all quantities on its later pages are invented. No historical committee approval, programme effect, realized saving or validated field power is claimed.

TECHNICAL ANALYSIS

Preserve the meaning before completing the row

[PD] The Row-to-Decision Evidence Contract (RDEC) links source identity, typed fields, same-version authority and joint package evidence. Each transition requires its own test.



Proposed DAG. Randomize Z across independent reviewer domains; do not adjust for later M or D.

[PD] Z is the assigned evidence workflow; M its source-linked representation; D the resulting decision; T actual implementation; Y the mature original-cohort outcome. C captures context and H real authority/dependencies. S supplies source versions and measurement. A change in M can improve a record without changing T or Y. H can constrain a decision even when its recorded label is wrong. The graph is proposed, not an empirically identified historical model.

[HR/AI] The register reports row reconstruction, role-field changes, metric normalization and open sizing questions. Some added intent was not filed on the board. Board assignee, submitter, lane coordinator and accepted accountable owner are different roles. Aligning a column to the assignee rule does not validate programme accountability. Named follow-up metrics already existed behind generic primary labels.

[AI] Preserve a typed value plus its evidence: estimate, confirmed explicit zero, documented unsized code, missing, unread or inapplicable. Coexisting numbers cannot authenticate a sentinel convention. The selected record contains conflicting descriptions of that convention. W3C tabular semantics and the Government Data Quality Framework support separating valid representation from accurate decision content; neither certifies authority or programme benefit.

Basis: proposed causal graph [PD]; critically read registry [HR/AI]; W3C (2015), Government Data Quality Hub (2020). No historical numerical coordinates are public.

TECHNICAL ANALYSIS

The same field totals permit opposite readiness

[AI] Independent hypothetical fixture: twelve original rows face three necessary, applicable gates. Each gate passes eight rows in both worlds.

Constructed register	Gate A	Gate B	Gate C	Jointly ready
Same eight rows pass every gate	8	8	8	8
Each gate fails a different four-row set	8	8	8	0

[AI] For N unique rows and k gates on one version, let n_j be each gate’s pass count. The sharp conjunction bounds are $\max(0, \sum(n_j) - (k-1)N) \leq \text{joint ready} \leq \min(n_j)$. Aligning failures attains the upper bound here; spreading failures across the entire universe attains the lower bound. Neither a weighted score nor the sum of filled cells selects a world. Even a measured intersection of populated fields may overstate valid readiness when the fields are stale, misassigned or unauthenticated.

Same standalone outcomes; common 100-unit cohort	Y(AB)	Joint benefit	Net AB
Same rescued units	70	30	5
Disjoint rescued units	40	60	35
Complementary response	10	90	65
Antagonistic combination	100	0	-25

[AI] Every world has $Y(\text{neither})=100$ and $Y(A)=Y(B)=70$ adverse outcomes. Net A is 20 after cost 10; AB costs 25. The preferred package reverses even though both standalone benefits are 30 and share the same population. The first three worlds never worsen any unit by adding an action; the last permits antagonism. Interaction on the benefit scale is joint benefit minus 60. These are explicit unit-level potential outcomes, not historical estimates or equally likely scenarios.

[PD] REVIEW INTERACTION; DO NOT AUTOMATICALLY DEDUPLICATE

Shared scope and equal estimates warrant a joint-effect test. They do not establish that the benefits should be summed or collapsed.

Basis: exact source-independent constructions [AI]; all gate sets exhaustively checked in a smaller universe. Knol and VanderWeele (2012) motivates common-reference interaction contrasts.

TECHNICAL ANALYSIS

Validate source fidelity and accepted authority separately

[PD] Freeze original request IDs, field definitions, source versions and the decision date. Keep filed, proposed, reviewed, approved and implemented states distinct.

Record unit	Required interpretation
Request × version	Original board identity, extraction time and later change history.
Person × role × request	Assignee, submitter and accountable owner, each with scoped acceptance.
Metric × population × horizon	Existing follow-up definition, numerator, denominator, maturity and direction.
Estimate × basis × authority	Gross/net value, currency or capacity units, costs, confidence and Finance scope.
Package × shared dependency	Common resources and population, interaction hypothesis and joint comparator.

[AI] Two hypothetical rows can retain their IDs, count and total value while their owners are swapped. A builder and verifier reading that same mistaken intermediate file will agree. Reproducibility therefore needs archived raw input, parser and intermediate versions, plus independent source-to-field comparisons and deliberate swap/null/unit tests. This is a logical vulnerability, not evidence that the historical parser made that particular error.

[PD] Reconcile every original row where feasible; otherwise sample with known inclusion probabilities across filed/nonfiled, assigned/unassigned, typed-value and shared-pool strata. Blind duplicate coders to workflow assignment and proposed funding. Report weighted confusion tables, prevalence, raw agreement and cluster-bootstrap kappa. Adjudicate against independently opened source panels and scoped owner/Finance acceptance, not a second rendering of the same extraction. An illustrative simple random sample of 400 gives worst-case binary agreement precision near ±4.9 percentage points, not clustered precision or kappa power.

[PD] Test invariance across business line, language, source version and field rule. Track missingness reasons and unresolved authority rather than deleting difficult rows. Reconcile row evidence at a fixed deadline and preserve later amendments. Wickham’s observational-unit discipline helps prevent one-to-many joins from turning role records or metric aliases into additional programmes.

Basis: proposed measurement protocol [PD]; independent semantic-swap test [AI]; Wickham (2014), Sandve et al. (2013), Cohen (1960).

TECHNICAL ANALYSIS

Test decision quality before claiming portfolio value

[PD] Begin with a bounded shadow benchmark. Controlled truth can test decision rules; it cannot make the historical joint programme effect observable.

[PD] Freeze original paired-request packets, common permitted evidence and authority conditions. Randomize independent reviewer domains to a current-evidence rule or RDEC. Retain simple one-package, both-package, unresolved and full-review policies; apply common blockers and evidence rights fairly. Lock choices at intake plus 72 hours and reveal an independent constructed reference at intake plus 21 days. In a field extension, reference work can establish source/authority fidelity, but causal joint value still needs an outcome design.

[PD] The benchmark estimand is reduction in expected decision regret per original packet, before additional review cost K on the same denominator. Separately estimate the unconditional probability of recommending a negative-net package, including unresolved and blocked packets. Charge review, delay and missed opportunity. Superiority to a risky comparator does not establish acceptable absolute harm; compare safe single-package choices too. A known authorized single package can remain feasible when joint evidence is unresolved; automatic deferral is a policy choice to compare, not a fact about unknown value.

[PD] A later approved operational factorial compares neither, A, B and AB on original eligible units in resource-sharing clusters. Estimate each assignment effect and the interaction $E[Y(AB) - Y(A) - Y(B) + Y(\text{neither})]$, with mature service outcomes, full costs and spillovers. Preserve necessary usual operations. A shadow review's lower regret does not prove a physical benefit. If rollout is staggered observationally, define group-time contrasts, pretrend equivalence and sensitivity; nonsignificant pretrends do not identify the effect.

[PD] Prespecify baseline evidence quality, accepted capacity, source age and package coupling as modifiers. Do not condition total ITT on later approval, completeness or implementation. Rows selected for poor quality invite regression to the mean. Under an assumed linear conditional-mean model, $E[\text{untreated follow-up} | x] = m + p(x - m)$, where x is selected baseline and m the population mean. Calibrate p and compare original cohorts; no historical value is estimated here.

[PD] Compare 80% MDE against K using domain-level variance and realistic recruitment. Compute finite null and boundary tails where tractable; normal planning alone is not calibrated coverage. More reviewers cannot repair a dominated decision rule, invalid reference or unidentified joint benefit.

Basis: proposed evaluation [PD]; Hudgens and Halloran (2008), Callaway and Sant'Anna (2021), Barnett et al. (2005). The compact numerical fixture is separate from these causal outcomes.

TECHNICAL ANALYSIS

Require evidence that could reject the interpretation

[PD] RDEC separates record reconstruction, readiness, authority, decision suitability and operating effect. Passing one test does not imply passing the others.

Proposition	Prespecified rejection or hold rule
P1: marginal completeness proves joint readiness	Reject using the exact counterexample; require same-row, same-version valid conjunctions.
P2: shared pools prove duplicated benefit	Reject the implication without joint effects; test alternative overlap and interaction mechanisms.
P3: RDEC adds economic decision value	Hold purchase unless cost and safety gates clear against relevant common-evidence alternatives.
P4: normalized fields transport across contexts	Hold extension when blinded error, acceptance or comparative value changes materially by prespecified context.

[AI] W3C and Wickham discipline field semantics and observational units; the Data Quality Framework distinguishes completeness from accuracy. Sandve supports executable lineage, while independent source checking addresses shared errors. Knol/VanderWeele requires separate component and joint contrasts. Deming cautions against tampering with sparse movements; Campbell/Goodhart predicts pressure to fill fields when completeness is rewarded. Ittner/Larcker tests the record-to-decision-to-outcome chain; Tucker/Edmondson redirects recurring placeholders toward process repair. Composite weights cannot compensate for absent critical authority.

[HR/RC] The frozen inventory contains 58 diagnostics: two sparse source endpoint comparisons, 25 versions, 14 design audits and 17 cross-sectional profiles. Neither source comparison crosses fitted X/MR limits, but neither establishes a stable operating phase. Laney is neither eligible nor computed. Tentative XmR constants are 2.66 and 3.267. Uncontrolled pilot endpoints do not identify intervention effects; no original row-level history is reconstructed from these aggregates.

Limitations and contribution

[HR/AI] Reported reconstruction, correction and evidence escalation support a provisional registry-level hoop. Exact originals remain unrecovered; a related parent planning workbook cannot authenticate the detail files. Metadata alone cannot establish sole authorship, sending, accepted accountability or approval. No customer benefit, implementation effect or review-productivity gain is established. Joint effects, prices, evidence errors, capacity and interference require calibration. The finite benchmark knows potential outcomes only by construction. Unknown sizing remains unknown; a short review-cycle result does not price indefinite delay.

Basis: proposed falsifiable contract [PD]; source boundaries [HR/RC/AI]; case-specific primary literature. This is an internal, agent-assisted methods assessment.

Test the combined recommendation while preserving a safe single choice

[PD] This wholly synthetic RDEC shadow replay tests a funding-packet recommendation. It changes no actual approval, spending, staffing or customer service.

[PD] Each original packet contains two proposals. Common evidence gives each standalone benefit 6 and action cost 4; the approved hypothetical budget is 8. Constructed joint benefit J is 6, 12 or 18 with probabilities 0.40, 0.40 and 0.20. Choosing one has net 2; choosing both has net J - 8. Common fields are identical across states, so the best current-evidence fixed choice is both.

Evidence law	Clear / ambiguous	Corroborated
Joint evidence resolves	0.90 / 0.65	0.98 / 0.90
Correct conditional on resolution	0.90 / 0.65	0.995 / 0.97
Clear-domain probability	0.70; stress 0.30	0.30

[PD] Wrong labels split equally between the other states. RDEC recommends one for a duplicate label and both otherwise. Unresolved joint evidence defaults to an independently authorized single proposal. A known filing/authority blocker affects 10% of packets: unsized is the only feasible action, so only delay is charged. All packets remain; confidence cannot waive a blocker.

[PD] Randomize independent review domains containing four packets, conditionally independent given a shared evidence grade. Distinct frozen batches have equal evidence rights. All assigned packets pay the fixed preparation/review cost, including blockers and unresolved cases.

CONSTRUCTED REFERENCE, NOT AN AUDIT ORACLE

[PD] Recommendations lock at original intake + 72 hours. The hidden benchmark J is released by intake + 21 days and never tunes decisions. A real file review cannot reveal all joint potential outcomes; empirical use needs a separate identification and reference-validation design.

[PD] Original unit is the two-proposal packet, not one row. Package costs are already inside net payoff. Common intake and research-reference budgets are shared, not free. The source supplies none of these probabilities or prices.

Size the decision comparison, not the number of completed fields

[PD] Primary ITT is decision-loss reduction per original packet. A second gate tests superiority in the unconditional probability of recommending a negative-net bundle.

Domains / arm	80% MDE	Price K	Zero-test power at BE	Full-law joint floor
Small / 16	Not reached*	0.35	17.4%	18.8%
Base / 128	0.336	0.35	83.2%	88.8%
Ambiguous / 128	0.365	0.35	76.8%	4.0%
Corroborated / 512	0.165	0.65	>99.9%	>99.9%

[PD] MDE and K use the same loss units per packet. Economic clearance requires the loss-reduction lower bound above K; the risk-reduction lower bound must exceed zero. Separate fixed normal radii use exact shared-grade moments and variance ceilings along the declared usual-to-RDEC mixture. Nominal directional error is 2.5%; these are approximate planning screens.

[PD] Base MDE 0.336 is below price 0.35; ambiguous MDE 0.365 is above it, despite positive mean net 0.02665. Corroboration changes evidence quality and costs, rather than treating a larger sample as a better policy. Risk superiority to both does not establish acceptable absolute risk; always-one has zero negative-net risk.

A CONDITIONAL DESIGN, NOT A RECRUITMENT PROMISE

[PD] The expanded proposal requires 4,096 packets across 1,024 domains; availability is unestablished. Selected finite zero-null loss errors are 2.28213% / 2.50410%, and risk errors 2.35081% / 2.44469%, for base / corroborated. Exact-law marginal tails give base full-law joint floor 88.7020%; they do not establish uniform coverage or justify purchase.

[PD] Four direct allocations, no sample-minimum search. *80% detection is not reached on the declared effect path. Eighteen finite marginal tails use polynomial convolution, not Monte Carlo; joint floors are bounds, not exact joint probabilities.

Price missed value, harmful recommendations and review together

[PD] Terminal loss is best feasible net value minus selected net value, plus 0.5 for an unsized response in this review cycle. Add fixed review K. These are decision judgments, not historical dollars.

Total loss / negative-net risk	Base	Ambiguous	Corroborated
Both / best fixed	1.4900 / 36.00%	1.4900 / 36.00%	1.4900 / 36.00%
One / zero negative-net risk	2.2100 / 0.00%	2.2100 / 0.00%	2.2100 / 0.00%
RDEC / safe fallback	1.1088 / 4.72%	1.4633 / 6.71%	0.9155 / 0.73%
Same-signal loss optimum	0.9828 / 11.03%	1.2654 / 16.61%	0.8608 / 3.47%
Fuller review	0.7465 / 0.36%	0.7465 / 0.36%	0.7465 / 0.36%

[PD] All four fixed actions and 256 same-signal mappings are checked per context. Always-unsized loses 4.4600. The signal optimum chooses both when unresolved: lower loss but greater risk. Base unresolved-hold loss 1.5025 exceeds safe-fallback 1.1088. Fuller review has resolution and accuracy 0.99, costs 0.65, and wins on loss. At price 1.25 its loss is 1.3465; base RDEC then costs less.

[PD] Base net versus both = $a \cdot h \cdot 1.251 - 0.51975 - 0.35$. Here a is realizability and h non-overlap of favorable category improvements. Split changes by state, grade and loss category first; added adverse losses and all review costs stay fully priced.

Net units / packet	h = 0.50	h = 0.75	h = 1.00
a = 0.25	-0.713	-0.635	-0.557
a = 0.50	-0.557	-0.401	-0.244
a = 1.00	-0.244	0.069	0.381

CHOOSE THE TRADE-OFF EXPLICITLY

[PD] Always-one has higher loss but zero negative-net risk. Fuller review has the lowest displayed loss, with residual risk. Agree prices and a risk criterion; passing against both does not select the best purchase.

[PD] Within declared loss support 0..10.5, uncertain reference on epsilon of original packets per arm permits contrast error $21 \cdot \epsilon$ for loss and $2 \cdot \epsilon$ for risk, with verified costs fixed. Other supports or eligibility require new bounds.

Buy evidence only when its decision value exceeds its price

[PD] This separate one-time research decision concerns a joint-benefit question, a separate-benefit question or neither. Its value is not added to recommendation gains or programme savings.

[PD] State probabilities are 0.50, 0.30 and 0.20. Choose a joint-effect study, a separate-benefit study or defer. Payoffs by state are (5,-1,-1), (-1,6,-1) and (0,0,0), including downstream action costs. Correct-signal probability changes by scenario; each error is equally split between the other states.

Evidence scenario	Correct signal	EVSI	Cost	Net EVSI
Low	40.0%	0.00000	0.40000	-0.40000
Base	70.0%	1.02000	1.00000	0.02000
High	90.0%	1.81000	1.80000	0.01000

[PD] All 27 signal-to-action rules are evaluated per scenario. No-information value is 2.00 and EVPI is 2.30. Weak information changes no optimal action and loses its entire price. The high-information scenario has lower net value than the base because it costs more. These are alternative probability and price judgments, not confidence intervals or historical findings.

START WITH A BOUNDED CURRENT-SOURCE CHECK

[PD] Recover the original row and panel evidence, resolve role and value conventions, and obtain one defensible joint-effect comparison. Calibrate query accuracy, blockers, review effort, prices and reference uncertainty before registering a replay. Retain a safe fallback and the option to defer research; obtain a committee ruling separately.

[PD] The shadow recommendation benchmark cannot establish Finance approval or physical programme incrementality. Full review and same-signal optima are conditional on the invented law and require external calibration before operational use.

REPRODUCIBILITY / READER PROTOCOL

Reproduce the row without upgrading its authority

[AI] The Row-to-Decision Evidence Contract separates successful reconstruction from the decision the record is meant to support. An analyst should be able to reproduce both the corrected artifact and every remaining uncertainty.

Reproduction layer	Required record and decision test
Original evidence	Identify the source version, row locator and extraction boundary. Retain inaccessible panels, unread fields and derivative summaries as distinct evidence states.
Transformation	Rebuild the stated schema and preserve observed estimates, placeholders, explicit zeros and absent values. Audit the written output, not only an intermediate object.
Meaning and acceptance	Track when the owner field changes meaning. Separate the recorded assignee, proposed accountable owner and accepted responsibility. Keep intent rows distinct from filed requests.
Joint economics	Reconstruct common cost pools and alternate programme values. A shared population triggers a joint-effect question; it does not supply a validated overlap deduction.
Decision evaluation	Rebuild the hypothetical state, signal, action and cost law. Include simple current-evidence rules and full review; retain unresolved requests and all acquisition costs.

[PD] For each claim, record the raw field, its interpreted meaning, the authority for that interpretation and the version in which it applies. An owner confirmation or funding ruling becomes a separate dated event. Reopening a record should reveal whether a substantive decision changed or whether the same evidence was merely displayed more clearly.

[AI] The private archive contains selected source rows, conflict arithmetic, state definitions, executable constructions, sensitivity checks and internal reviews. It cannot supply an original export that was not recovered, turn a reported script description into observed execution, or convert request totals into Finance-approved benefits. The public examples use independent invented values, not concealed historical rescaling.

Protected evidence S1 and current recovery audit; independent calculations; proposed RDEC reproduction protocol. Historical source coordinates remain private.

REFERENCES / 1

Primary sources and their role in this case

[AI] The cited methods clarify provenance, measurement and decision evaluation. They do not authenticate the historical source or validate the programme requests carried by the register.

NIST/SEMATECH. Individuals Control Charts, e-Handbook section 6.3.2.2. Assess actual ordered observations before computing process limits. Request categories, version pairs and proposed targets do not form a stable operating time series.

Laney, D. B. (2002). Improved Control Charts for Attributes. *Quality Engineering*, 14(4), 531–537. Require original event counts and rational subgroups. Completeness percentages from one portfolio snapshot cannot establish a changing process.

Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20, 37–46. Audit independently coded field states and owner roles. Separate agreement with a codebook from the correctness of the underlying claim.

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. A deliberately selected group of problematic requests can become less extreme without an effective reconciliation treatment. Freeze selection before evaluation.

Callaway, B., & Sant’Anna, P. H. C. (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, 225, 200–230. For later phased use of a reconciliation process, define adoption cohorts and explicit comparison domains; assess anticipation and changing intake mix.

Collier, D. (2011). Understanding Process Tracing. *PS: Political Science & Politics*, 44, 823–830. Attribute the reconstruction, interpretation rules and exception disclosure separately from source ownership, funding authority and programme outcomes.

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. Value evidence through the decision it changes under the same feasible actions. Subtract acquisition costs and avoid adding information value to the same claimed operating benefit.

Primary author, journal, institutional and official-method sources; access scope and case-specific applications are recorded privately.

REFERENCES / 2

Primary sources and their role in this case

[AI] The cited methods clarify provenance, measurement and decision evaluation. They do not authenticate the historical source or validate the programme requests carried by the register.

The W. Edwards Deming Institute. The Funnel Experiment. Repeatedly rewriting a register in response to sparse status changes can create management noise. First establish comparable observations and decision rules.

Campbell, D. T. (1976). Assessing the Impact of Planned Social Change. A completeness target can reward filled cells that erase uncertainty. Preserve unread, unsized and unresolved states even when the register appears less complete.

Goodhart, C. A. E. (1992). The Objectives for, and Conduct of, Monetary Policy in the 1990s. Reserve Bank of Australia Annual Conference, section 1(b)(i). A scored planning field can become detached from decision quality. Do not allow a higher confidence label to substitute for approval or benefit validation.

Ittner, C. D., & Larcker, D. F. (2003). Coming Up Short on Nonfinancial Performance Measurement. Harvard Business Review, 81(11), 88–95, 139. Test the chain from a better register to a better decision and then to mature operating outcomes. Formatting success does not identify either downstream effect.

Tucker, A. L., & Edmondson, A. C. (2003). Why Hospitals Don't Learn from Failures: Organizational and Psychological Dynamics That Inhibit System Change. California Management Review, 45(2), 55–72. A one-time reconstruction can conceal recurring failures at intake. Record unresolved source defects and test whether the accountable source owners close them.

OECD / European Commission JRC (2008). Handbook on Constructing Composite Indicators: Methodology and User Guide. Separate field coverage, source traceability, accepted authority and economic evidence. A weighted confidence score must not compensate for a failed mandatory condition.

W3C CSV on the Web Working Group (2015). Model for Tabular Data and Metadata on the Web. W3C Recommendation, 17 December 2015, sections 4.4–4.6 and 6.4. Preserve row keys, field types, null meanings and original source positions before checking completeness.

Primary author, journal, institutional and official-method sources; access scope and case-specific applications are recorded privately.

REFERENCES / 3

Primary sources and their role in this case

[AI] The cited methods clarify provenance, measurement and decision evaluation. They do not authenticate the historical source or validate the programme requests carried by the register.

Wickham, H. (2014). Tidy Data. *Journal of Statistical Software*, 59(10), 1–23. doi:10.18637/jss.v059.i10. Separate request, person-role, metric definition and approval-event units; join them with explicit keys.

Government Data Quality Hub (2020). The Government Data Quality Framework. Published 3 December 2020, core data quality dimensions. Report completeness, validity, consistency and accuracy separately; scope quality to the funding decision.

Sandve, G. K., Nekrutenko, A., Taylor, J., and Hovig, E. (2013). Ten Simple Rules for Reproducible Computational Research. *PLOS Computational Biology*, 9(10), e1003285. doi:10.1371/journal.pcbi.1003285. Archive input, parser, intermediate data and verification lineage; a repeatable build still needs an independent source check.

Knol, M. J., and VanderWeele, T. J. (2012). Recommendations for Presenting Analyses of Effect Modification and Interaction. *International Journal of Epidemiology*, 41(2), 514–520. doi:10.1093/ije/dyr218. Display each component and their joint contrast against the same reference before attributing duplicated or complementary benefit.

W3C; Moreau, L., and Missier, P. (eds.) (2013). PROV-DM: The PROV Data Model. W3C Recommendation, 30 April. Sections 2.1 and 4.2–4.3 distinguish a versioned entity, its generating activity and responsible agent, including the source of a provenance assertion. Application: retain the source field, transformation and decision record separately. Provenance structure does not authenticate a funding decision or establish a causal benefit.

Hudgens, M. G., & Halloran, M. E. (2008). Toward Causal Inference with Interference. *Journal of the American Statistical Association*, 103, 832–842. Original manuscript. Randomize shared reviewer and portfolio domains because standards, dependencies and review resources can affect multiple requests.

Primary author, journal, institutional and official-method sources; access scope and case-specific applications are recorded privately.

FINAL REFEREE REPORT / INTERNAL REVIEW

Publish the reconstruction; preserve the unresolved decisions

[AI] Author-led, agent-assisted review; not external peer review. The case contributes a reproducible distinction between field completeness, valid interpretation, authorized decisions and operating effects.

Major issue	Resolution and decision boundary
Coverage confused with accuracy	Separate a filled field from accepted meaning. Preserve conflicting versions, unread panels and source-derived versus proposed owner assignments.
Normalization described as new capability	A named metric recovered from a follow-up answer is extraction and normalization. A changed owner convention is record alignment, not a validated accountability error rate.
Overlap deducted without identification	Common populations and identical figures flag a reconciliation need. Retain alternative joint effects; neither mechanical summation nor mechanical deduction is justified without the relevant evidence.
A polished row treated as approval	Keep authored confidence, filed status, accepted funding and mature results distinct. The record does not establish downstream approval or realized savings.
Model treated as deployment evidence	Restrict each benchmark to feasible actions. Retain mandatory holds, all review costs, the safe single-package option and stronger full review. Hypothetical precision is not deployment evidence.

Partner critique: can a COO act on page one?

[AI] Yes. Accept the reconstructed register as a review instrument, assign unresolved source questions to their accountable owners, and require a scoped decision record before promoting any row to a commitment. Ask for joint effects and incremental costs when requests share a cost pool.

Remaining limits and minor issues

[AI] Source authentication, field semantics, accepted ownership, overlap and funding closure constrain the conclusions. The technical section consolidates these limits. No internal calculation, proposed proof or reviewer agreement converts a reported build into demonstrated business impact.

REVIEW DISPOSITION / PASS WITHIN STATED SCOPE

[AI] Publish the bounded reconstruction and testable decision design. Retain the distinction between a more readable portfolio and a portfolio whose claims have been approved or delivered.

Final source, numerical, methodological, partner and actual-page reviews are required. This is internal publication review, not external assurance or endorsement.