

OPERATING CASE STUDY / STRATEGY & CAPITAL ALLOCATION

Enterprise Operating Plan Design

Bind each approval to the plan it actually reviewed

An operating-project case on annual planning in a distributed service network. It examines portfolio reconciliation, resource commitments, benefit validation and the evidence required to move from a submitted workbook to an executable plan.

PORTFOLIO PLANNING & EXECUTION READINESS

BLUF / BOTTOM LINE UP FRONT

Accept the submission for review; authorize execution only against a verified plan version.

ANALYTICAL FINDING

A balanced staffing total does not establish a feasible schedule. A populated workbook does not establish accepted ownership. Later edits can invalidate specific approvals even when the earlier submission was complete.

CONTRIBUTION IN THE SOURCE RECORD

Registry-level hoop: reported plan authoring, reconciliation and escalation. Funding, owner acceptance and execution belonged to other roles.

EVIDENCE STATUS

Evidence-reconciliation and operating-design case. Historical originals remain unauthenticated; execution and realized benefits are unverified.

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

STRUCTURED ABSTRACT

Separate readiness to submit from authority to execute

[AI] The useful planning output is an explicit decision contract: what is proposed, what is funded, what remains unresolved, and who can authorize the next stage.

Element	Decision-relevant conclusion
Context [HR]	A planning lead assembled an annual operating portfolio from narrative requirements, project rows, staffing calculations, performance goals and benefit estimates held in different records.
Objective [HR/AI]	Produce an intelligible submission while exposing unresolved capacity, ownership, measurement and financial assumptions. Avoid converting spreadsheet completeness into implied approval.
Approach [RC/AI]	Reconcile versions, row populations, value bases and resource definitions. Test whether each approval still applies to the current portfolio and whether annual allocations can support the delivery sequence.
Findings [AI]	A funded total and an effort estimate answer different questions. A known subtotal is not a complete valuation. Equal completeness scores can conceal different critical blockers.
Decision value [PD]	The planning sponsor should accept an interim submission with named issues, then withhold execution lock until indispensable approvals match the current version and scope.
Implications [PD]	Test the proposed Versioned Portfolio Evidence Contract (VPEC) against snapshot reliance and full rereview. Evaluate decision errors and review cost before claiming operating benefit.

[PD] Finance validates value basis and incrementality; resource owners confirm deployable capacity; operating owners accept actions and dates. The planning lead reconciles the package and routes unresolved decisions. No one role can substitute its spreadsheet edit for all three approvals.

[HR] historical record; [RC] recalculated; [AI] analytical inference; [PD] proposed design. Protected evidence S1; original recovery and claim ledger retained privately.

EXHIBIT 1 / DECISION ARCHITECTURE

A submission can be complete while execution remains unauthorized

[AI] A single green status merges three distinct decisions. Preserve the conditions and evidence needed for each transition.

Stage	What the decision permits	Required evidence
Submit	Enter the formal review process	Required fields, traceable assumptions, version identifier and explicit unresolved issues.
Interim acceptance	Continue planning within stated constraints	Named issue owners and dates; a permitted review scope; no implied release of blocked resources or actions.
Execution lock	Commit the approved operating package	Current scope, accepted ownership, feasible resources and schedule, and all indispensable approvals required by policy.

[HR] The record describes progressive completion of the submission, correction of manager comments and escalation of open questions. It also contains readiness checklists captured at different moments. An early incomplete checklist cannot establish that the final submission failed; a later populated row cannot establish that all external approvals were obtained.

[PD] State a gate as a conjunction of required conditions, not an average score. Record a deliberate exception as an exception with authority, scope and expiry. A blocker can be appropriate for execution lock while still allowing a useful interim submission.

EXECUTIVE RULE

[PD] Ask which decision the evidence authorizes. Do not ask only whether the workbook is finished.

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

EXHIBIT 2 / CHANGE CONTROL

Revalidate the approvals affected by a material plan change

[AI] A valid review is attached to a version and scope. It does not automatically approve every later edit, nor does every edit require a complete restart.

Change	Approval dependency to examine	What cannot be inferred
Additional project or resource demand	Capacity, prioritization and schedule	That a changed total is physically staffed or unauthorized.
Changed benefit estimate or unit cost	Finance basis and non-overlap	That a larger number represents incremental, avoidable value.
Different assigned owner	Acceptance and delivery responsibility	That assignment proves acceptance, authorship or funding authority.
Editorial correction only	Meaning and evidence locator	That every wording change invalidates unrelated approvals.

[HR/AI] The source records a balanced submission and a later shared version with additional rows and staffing. The bridge between these versions is arithmetically explainable. That reconciliation does not identify who authorized the change, whether the revised scope was accepted, or which version became the execution baseline.

[PD] Preserve a semantic change log linking fields to decision dependencies. Reopen only affected conditions when that mapping is reliable; send ambiguous changes to full review. Match the reviewed version to the released version atomically with lock, so a concurrent edit cannot bypass the check. A correct earlier review cannot protect an unchecked later snapshot.

[AI] This is configuration control applied to an operating portfolio. It preserves legitimate coauthoring while making approval scope observable. The proposed gate is useful only if the change log and dependency map are themselves complete.

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

EXHIBIT 3 / CAPACITY AND SEQUENCING

An annual staffing balance cannot certify a delivery schedule

[AI] Funded allocation, estimated work and deployable capacity are separate quantities. Even equal annual totals can hide a timing constraint.

Quantity	Definition	Decision implication
Funded allocation	Authorized resource envelope assigned to the portfolio	Constrains expenditure or headcount; does not establish that named capacity is available at every required time.
Estimated demand	Task effort under stated frequency and productivity assumptions	Identifies the work implied by scope; retain recurring and one-time work separately.
Schedulable capacity	Skill-specific availability across dependencies and deadlines	Determines whether accepted work can finish when promised. Annual aggregation loses this information.

[HR/AI] The plan allocates the funded envelope while the effort model indicates a larger demand. Deferring selected lines changes scope; proportionally reducing allocations rations coverage without reducing the work estimate. Neither operation establishes that remaining commitments fit the people, sequence and governance workload.




EXHIBIT 4 / VALUE RECONCILIATION

Resolve the value basis before choosing a headline total

[AI] Alternative totals are often different objects, not competing estimates of the same object. Treat their variation as a reconciliation problem before treating it as uncertainty.

Difference	Required reconciliation	Do not do this
Different project populations	Stable project key; include, exclude, combine and defer decisions	Average totals from unequal portfolios.
Unknown amounts or sentinels	Typed missingness and explicit sizing status	Convert blanks or a placeholder token into zero economic value.
Different benefit horizons	Annual run-rate, within-year realization and cash timing	Compare an annual opportunity with a partial-year receipt.
Shared outcome pool	Unit-level overlap and incremental mechanism	Assume all same-pool projects duplicate each other, or all are additive.

[HR/RC] The source contains several plausible portfolio totals built from different line items, scope cuts and maturity claims. A manager-level value does not equal the corresponding whole-line sum. No whole-line subset reproduces that manager figure; partial scope, a haircut or another basis could still explain it if documented. The disagreement must remain unresolved until that basis is recovered.

[PD] Maintain a reconciliation bridge from each version to the selected decision basis. Publish a known subtotal with an explicit incomplete-population label when some rows remain unsized. Keep typed placeholders out of arithmetic. A numeric cell containing a status sentinel is not an estimated benefit.

[AI] The case supports disciplined refusal to publish an unsupported grand total. It does not prove that the excluded opportunities lacked value or that a lower figure was the correct economic forecast.

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

EXHIBIT 5 / FINANCIAL EVIDENCE

Reproducible arithmetic is one gate below validated value

[AI] An analyst can reproduce a benefit formula without establishing that the effect is incremental, avoidable, mature or attributable to the proposed work.

Evidence level	What it establishes	Next question
Formula reconstructed	Inputs and arithmetic can be recomputed	Are units, population, horizon and source versions valid?
Mechanism supported	The proposed work can plausibly affect the outcome	What would happen without this package, and who else acts on the same units?
Finance validated	An accountable reviewer accepts the stated valuation basis	Does the approval apply to this scope and version?
Benefit realized	A mature outcome or cash record is observed	Is the change attributable and incremental after costs?

[HR/AI] The record distinguishes a small reconstructible subset from a broader set of unvalidated estimates. It also contains conflicting claims about financial sign-off. Under the conservative publication rule, no authenticated approval is promoted from that conflict. This is an evidence-status conclusion, not an estimate that the portfolio has no economic potential.

[PD] Tie the economic claim to an outcome estimand: eligible mature exposure × incremental effect × avoidable unit cost × non-overlap share, less incremental implementation and operating cost. Validate each factor separately. A conceded payment, an accounting allocation and an avoidable cash expense cannot share one unit-cost label.

[AI] A previous expected benefit and a later partial-year realization do not form a clean success rate without aligned scope and maturity. Use such records to locate assumption risk; do not attribute their shortfall to the planning lead or infer this portfolio’s future effect.

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

EXHIBIT 6 / MEASUREMENT CONTRACT

A target needs a stable population and a feasible action bridge

[AI] A target can be arithmetically correct and operationally unsupported. The path from baseline to target must name its exposure, measurement convention and funded actions.

Contract component	Required test	Failure interpretation
Metric identity	Distinguish aliases, technical identifiers and business labels	A filled name is not verified lineage.
Population and units	Match numerator coverage, denominator and business line	A recalculated ratio may still measure the wrong construct.
Baseline and maturity	Align year-to-date, weekly and annual horizons	A version change is not a process improvement.
Action bridge	Sum incremental action effects under shared capacity and overlap constraints	A retained top-down goal is not a funded bottom-up forecast.

[RC/AI] The lineage categories in the source account for fewer metrics than its stated total. That residual is a reconciliation question, not proof that the remaining metric is invalid. Similarly, technical-name gaps limit traceability even where every visible business-label field has been populated.

[HR/AI] Several goals retain fixed targets while the documented action bridge leaves a gap. Other metrics lack a verified baseline or definition in the accessible evidence. Preserve these as distinct states: unverified lineage, incomplete bridge and unsized metric. The corrective owner and required evidence differ for each.

[PD] Use one metric contract per eligible population. A source version may supersede a preliminary baseline; retain the earlier value as history rather than averaging them. Require a bottom-up sensitivity to volume, timing and shared actions before calling the target executable.

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

EXHIBIT 7 / DECISION REGISTER

Close the critical conditions, not the largest number of cells

[AI] Completion counts summarize paperwork. Execution readiness depends on which conditions remain unresolved and who has authority to close them.

Decision	Accountable role	Evidence needed to close
Which portfolio is in scope?	Planning sponsor	Versioned row set, deferrals and documented scope changes.
Can the work be delivered?	Resource and operating owners	Accepted actions, skill/time capacity and dependency schedule.
Which value is admissible?	Finance owner	Population, horizon, incrementality, cost and overlap decisions.
Which goal is measurable?	Measurement owner	Metric lineage, denominator, maturity and baseline version.
May the package lock?	Authorized sponsor	Current gate evidence, explicit exceptions and final version match.

[PD/RC] Independent illustration: two plans each pass nine of ten checks. The first fails indispensable resource authorization; the second lacks an optional explanatory note. Both score 90%, but only the second can lock under the declared rule. The relevant statistic is satisfaction of critical conditions, supplemented by the unresolved-issue register.

[PD] An unresolved item should carry its decision, evidence request, owner, deadline and effect on submit/interim/lock status. Aging should trigger a choice: resolve, defer scope, accept an authorized exception, or stop the affected commitment. Repeatedly moving the deadline is not closure.

[AI] The planning contribution is the creation of an auditable decision package and the exposure of uncertainty. Historical execution, causal service improvement and realized savings require separate evidence. The prospective evaluation below tests the gate’s decision quality within this narrower scope.


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

TECHNICAL ANALYSIS

Submission and final authority are different states

[PD] The Versioned Portfolio Evidence Contract (VPEC) binds each gate to a portfolio version, decision scope and evidence date. Interim submission can preserve momentum while named dependencies remain unresolved.

VPEC: A REVIEW CHANGES EVIDENCE AND MAY OBTAIN NEW AUTHORITY



Z also changes review work included in L. C can affect review, decision and deadline validity.

K is not programme economic truth. O must reconstruct K independently; ascertainment may err.

Proposed shadow DAG: no observed customer-benefit path or historical causal coefficient is asserted.

[PD] Randomized policy Z changes impact review and authorized reaffirmation R, the lock/interim decision D and review work. Context C includes original scope, evidence quality, capacity and edit history. K is the actual version/authority state at the decision deadline; O is its independently reconstructed reference. Decision-resource loss L depends on the decision, validity and work. Arrows are proposed mechanisms, not estimated historical effects.

[AI] Reaffirmation may legitimately change K; therefore the benchmark evaluates the whole review policy, including its authority response, rather than classifying every case against an immutable pre-review label. Missing approval cannot authorize a lock. A final atomic compare-and-lock protects against a tracked in-file race. Valid final authority also requires the signed dependency manifest to cover relevant external states.

[HR/AI] The registry reports assembled submission artifacts, source reconciliation and escalated decisions. It does not authenticate final central acceptance, Finance validation, delivered programmes or savings. An assigned row owner does not identify the editor of a later change. A changed shared-file total requires version and authority reconciliation; it does not, by itself, prove unauthorized conduct.

Basis: proposed VPEC causal structure [PD]; source stage and role boundaries [HR/AI]. NIST SP 800-128, section 3.3; GovS 002, clauses 7.7-7.9, used as methodological analogies.

TECHNICAL ANALYSIS

A high readiness score cannot authorize a lock

[AI] Exact independent hypothetical constructions. No values below reproduce historical planning quantities. A score summarizes progress; the declared critical gates determine permission.

Same reported score	Unmet item	Valid final lock?
Nine of ten checks pass	One mandatory authority gate	No
Nine of ten checks pass	One explicitly optional note	Yes, if all other declared conditions hold

[AI] Let g_j be verified gate status and C the declared critical set. Lock requires AND over j in C of $g_j = 1$, together with a current version match. With nine mandatory checks and one optional check, all ten configurations with exactly one failure score 90%; nine fail the mandatory conjunction. Changing weights cannot replace that rule unless the governance authority explicitly changes what is mandatory.

[AI] Funding also has a timing constraint. Two projects each need one specialist for one full period before the end of period 1. One specialist is available for two periods. Total demand of two specialist-periods fits the annual capacity of two, yet no schedule meets both deadlines. Moving one deadline to period 2 permits one feasible order. This is a logical counterexample, not a diagnosis of the historical capacity model.

[PD] Record the dependency edges and invalidate only affected approvals. A scope edit can propagate through capacity, schedule and Finance gates; a genuinely nonmaterial title edit need not. Revalidation must use the authorized materiality rule, not the author’s convenience. Preserve the prior approval and its scope rather than overwriting history.

[AI] NEW DECISION RULE

Require the critical-gate conjunction, a feasible dependency schedule and version-bound authority. Neither an exact funded total nor a high completion percentage substitutes for those conditions.

Basis: exhaustive independent gate and two-period schedule enumeration [AI]. Full proof and input hashes are private. GAO Schedule Assessment Guide provides the scheduling context.

TECHNICAL ANALYSIS

Reconstruct the portfolio without manufacturing value

[PD] An arithmetic total, an approved economic estimate, a funded allocation and an executable plan answer different questions. Preserve those statuses in separate fields and separate roll-ups.

Object	Evidence required before aggregation
Project benefit	Common counterfactual, eligible exposure, outcome and cost grain, currency/date, maturity and incremental effect.
Portfolio benefit	Joint treatment scope, overlap and interaction; net adverse effects; complete incremental implementation and operating cost.
Approval	Named authorized decision, applicable scope and version, conditions and expiry; a proposed assignee is not acceptance.
Unknown or unsized	Typed unresolved status with source owner and due date; exclude from a validated subtotal without declaring economic value zero.

[AI] Shared outcome populations identify a need to reconcile overlap, not proof that every project dollar is duplicated. Conversely, adding internally coherent lines cannot establish the joint causal benefit. For a frozen portfolio S, value is $V(S) = \sum_t \text{discount}_t \times \{E[\text{cost}_t(\text{usual})] - E[\text{cost}_t(S)]\} - \text{incremental implementation cost}$. The same cost cannot also be embedded in the outcome term.

[AI] Even identical marginal forecasts leave joint risk open. Suppose each of two projects yields either 0 or 10 with equal probability. If their outcomes move together, total value is 0 or 20; if they move oppositely, it is always 10. Both have expected total 10 and the same individual ranges. The probability of falling below 10 differs from one-half to zero. These invented distributions illustrate dependence; the source supplies no such probabilities.

[HR/AI] Conflicting header scopes, alternative goal baselines and pre-fill versus later checklists remain typed versions. An allowed identifier can still reference the wrong construct. Year-to-date does not automatically resolve maturity or numerator scope. A recovered planning artifact has unresolved provenance and version; it does not authenticate the registered submission. A bounded search cannot prove approval never existed.

Basis: portfolio valuation identity and joint-risk construction [AI]; registry boundaries [HR/AI]. Phillips and Bana e Costa (2005); GAO-20-195G; HM Treasury (2026).

TECHNICAL ANALYSIS

Validate the change decision before testing delivery

[PD] Start with a shadow benchmark of original portfolio submissions. It measures version/authority decisions and work; it cannot establish the economic truth or eventual operational success of those portfolios.

[PD] Freeze the original submission and approval scope before reviewer-domain randomization. Compare snapshot reliance, targeted dependency review with a final version check, full rereview with the same final check, and interim submission. Give equal evidence and authority access. At intake + 48 hours lock the shadow response; by +7 days a blinded reference reconstructs the original-deadline state. Later approval cannot retrospectively make an earlier unsupported lock correct.

[PD] Define $\text{delta_L} = E[L(\text{snapshot}) - L(\text{VPEC})]$ per original portfolio, with every blocker and interim case retained. Charge incremental review/setup k on that same denominator; economic clearance requires a lower bound above k . Separately require fewer wrong locks per original portfolio. Repeated portfolios sharing a reviewer, dependency or authority queue belong to one cluster; do not select only approved or reaffirmed cases.

[PD] Validate measurement with a stratified probability sample of safe edits, material changes, missing responses and disputed versions. Independent coders blinded to assignment use original approvals and edit logs, then an authorized adjudicator resolves disagreements. Report weighted confusion tables, prevalence, agreement and cluster-bootstrap kappa. A hypothetical simple random sample of 400 limits ordinary agreement precision to about ± 4.9 percentage points at 95%; it is not a kappa-power calculation or valid after ignoring clustering.

[PD] Estimate cluster covariance, reference errors and review-delay distributions in a feasibility phase. A rough equal-arm screen is $\text{MDE}_{80} \approx 2.80 \sqrt{2 \sigma_c^2 / m}$, for m independent domains per arm and cluster-summary variance σ_c^2 . Compare with k in decision-loss units; calibrate the actual estimator at null, economic and safety boundaries. Conditional invented laws later in this case are sensitivity exercises, not field power. More samples cannot overturn a cheaper interim or full-review comparator.

[PD] A later field trial needs separate original operating cohorts, randomized rollout by resource-sharing domain, delivered-action fidelity, independent mature customer outcomes and full operating costs. A valid shadow approval does not identify that field ITT. Preserve spillovers, delay and negative effects rather than counting withheld expenditure as realized customer benefit.

Basis: proposed measurement, estimands and design [PD]. Cohen (1960); Hudgens and Halloran (2008). No historical matched panel, deployment effect or achieved power is asserted.

TECHNICAL ANALYSIS

Specify the observations that could reject VPEC

[PD] The framework is untested. A funding gate that remains imprecise requires a hold; it is not proof of no effect. Predeclare the meaningful thresholds and strong comparators.

Proposition	Disconfirming evidence or hold rule
P1: critical gates protect authority	Reject if a version with a failed mandatory gate is still certified under the specified rule; an optional omission alone must not cause rejection.
P2: scoped revalidation reduces stale locks	Require fewer wrong locks on all original submissions. Prespecified failure to clear safety or cost thresholds blocks adoption.
P3: targeted review conserves resources	Compare total loss with full rereview and interim submission in the same evidence context. Superiority over a weak snapshot rule is insufficient.
P4: readiness supports executable delivery	Audit dated specialist capacity, accepted responsibilities and dependencies; reject a claimed feasible schedule when no permitted allocation meets the deadlines.

[AI/PD] Stratify by preassignment edit complexity, external-dependency burden and authority-response capacity. Realized approvals and late mutations are postassignment outcomes, not primary-effect adjustment variables. Under an assumed conditional-mean model, $E[\text{untreated follow-up} \mid x] = m_0 + \rho \times (x - m_0)$, where x is the selected baseline and m_0 the population mean. Assumed $\rho = .4/.7/.9$ implies 60%/30%/10% reversion of selected excess. These hypothetical sensitivities quantify selection risk, not estimated fractions.

[AI] NIST and GovS 002 motivate controlled versions, not automatic truth. Phillips and Bana e Costa distinguish deliberation from accountable allocation; GAO separates estimate traceability from risk; Green Book preserves unpriced effects. Deming cautions against tampering with sparse signals; Campbell/Goodhart against rewarding filled cells; Ittner–Larcker requires a tested measure-to-value chain; Tucker–Edmondson highlights workarounds; composite-indicator weights cannot conceal mandatory gates.

Contribution and limitations

[RC] The private inventory has 5 source/sensitivity screens, 23 versions, 16 designs and 15 profiles. Tentative XmR diagnostics have no source fitted X/MR crossings; no stable phase is established. No eligible Laney panel is supplied, so Laney is not computed. Sparse annual-to-partial-year pairs do not become a stable operating series.

[HR/AI] The reported assembly, standard-setting and escalation support a provisional registry-level hoop. Originals, exclusive authorship, central approval, operating adoption and realized benefit remain unauthenticated. No complete joined outcome panel or valid historical comparator identifies impact. Shared-file authority, external dependencies, resource schedules, valuation overlap and reference accuracy remain material assumptions. The exact constructions change the release rule; they do not show which historical failure occurred.

Basis: proposed falsification and moderation [PD]; research synthesis and bounded contribution [AI/HR]. SPC eligibility and exact private inventory are documented in the separate evidence audit.

Approve the version that will actually be used

[PD] The Versioned Portfolio Evidence Contract (VPEC) separates a permitted interim submission from a final lock. A lawful edit may invalidate an earlier scoped approval without being an unauthorized act.

[PD] Randomize independent review domains, each with 12 original portfolio replays. Each starts with a valid scoped approval. Decisions lock at intake + 48 hours; a hidden reference by +7 days reconstructs the version and authority at that deadline. This is shadow work: no funding, customer promise or live permission changes.

Synthetic early edit	Share among unblocked	Targeted sensitivity: clear / ambiguous
safe edit	60%	8% / 18% false flags
capacity dependency	18%	95% / 65%
value-basis dependency	14%	80% / 45%
owner dependency	8%	90% / 60%

[PD] A common 10% known-blocked group stays interim. An impact flag requests authorized reaffirmation, arriving with probability 70% / 40% by grade. Later material in-file edits occur with probability 8% / 20%. Both reviewed rules atomically compare-and-lock the final version: a late change keeps the submission interim. Hidden edit class is unavailable to the policy; approval does not prove economic truth.

A version match is necessary, not sufficient

[PD] The model assumes the tracked dependency graph is complete. A matching file hash cannot validate an omitted external dependency, programme benefit or executable schedule. Unvalidated value remains unknown, not zero.

[PD] Entire law is invented. Base clear/ambiguous mix is 75/25; stress is 25/75. Atomic version check and timely independent reference are testable idealizations.

Test precision at the price of the safeguard

[PD] Primary ITT is reduced terminal decision loss per original portfolio; break-even is 0.65 resource units. A separate gate requires fewer wrong locks across every assigned portfolio, including interim submissions.

Domains per arm / context	80% MDE	Zero-test power at break-even	Joint floor at full law
Small base / 8	1.971	14.2%	99.2%
Base / 64	0.709	72.8%	>99.9%
Ambiguous / 64	0.715	72.0%	>99.9%
Expanded ambiguous / 128	0.506	95.0%	>99.9%

[PD] Exact event moments include a shared evidence-quality grade within each domain. Fixed variance-envelope radii give approximate one-sided 97.5% bounds. Investment needs the loss-reduction lower bound above 0.65, plus the wrong-lock lower bound above zero. The joint floor combines approximate marginal powers; it is not exact coverage.

[PD] Both 64-domain candidates remain below 80% power at break-even. The expanded design uses 3,072 portfolios and reaches 95.0% conditionally. MDE tests a positive effect against zero; investment needs clearance above 0.65. Per-portfolio rule mixing is not observed adherence; no minimum or recruitment feasibility is claimed.

The base design is not confirmatory-ready

[PD] The initial 2,500 base-zero trials gave 3.32% positive rejection (95% interval 2.69-4.10%). Independent finite-law convolution gives 2.49756% for loss and 2.52006% for wrong locks: the first run does not establish systematic inflation. Base 64 remains underpowered. Expanded sampling is a conditional screen; preregister and independently validate inference before confirmation.

[PD] Selected finite-process zero, economic-boundary and full-law checks report Wilson intervals. They do not establish performance over all unknown field distributions.

PROPOSED VERSION TEST / 3 OF 4

Keep the inexpensive interim option in the decision

[PD] Hypothetical penalties are 10 for a wrong lock, 2 for an unnecessarily interim ready version and 0.5 for an interim unready version. These are decision-resource prices, not programme cash or lost entitlement.

Policy	Base total	Ambiguous total	Extra work
Snapshot reliance	4.344	4.668	0.00
Targeted + version check	1.525	2.022	0.65
Full rereview + version check	1.646	1.814	1.25
All interim	1.321	1.272	0.00

[PD] All-interim is cheaper in both contexts at these prices; full rereview also beats targeted review under ambiguity. Common intake costs 0.10. Targeted reserved review/owner-response work, final check and setup cost 0.45 + 0.05 + 0.15 on every portfolio, including blockers. Full rereview costs 1.20 + 0.05. Delay prices 1, 2 and 4 preserve every comparator; interim does not set possible value to zero.

[PD] $Net = a \cdot h \cdot 3.605 - 0.136 - 0.65$. Here a is realizable decision-resource value and h is the non-overlap share. Favorable category reductions alone are discounted; additional interim losses and work remain fully charged.

Net units / portfolio	$h = 0.60$	$h = 0.80$	$h = 1.00$
$a = 0.50$	0.295	0.656	1.016
$a = 0.75$	0.836	1.377	1.918
$a = 1.00$	1.377	2.098	2.819

Reference and boundary limits

[PD] Unknown or wrong labels on epsilon of each original arm permit contrast errors up to 20*epsilon loss units and 2*epsilon wrong-lock units. External dependency omissions need a separate audit; the file-version guard does not cover them.

[PD] One review cycle only; repeated deferral needs repricing. Fix prices independently before confirmation. Research reference collection is separately budgeted; no historical avoided loss is asserted.

Use information to choose the unresolved decision

[PD] A separate one-time research choice concerns a capacity dependency, a benefit-basis dependency or no actionable dependency. Information value is not added to version-control savings.

[PD] The hypothetical prior is 0.35 / 0.30 / 0.35. Choose schedule reconciliation, benefit-basis audit or defer research. Payoff vectors are (6,0,-1), (0,7,-1), (0,0,0); they include downstream decision costs. The evidence-acquisition charge below is additional.

Evidence scenario	Correct signal	EVSI	Cost	Net EVSI
Low	50%	0.52500	0.50000	0.02500
Base	70%	1.15500	1.00000	0.15500
High	90%	1.99500	2.00000	-0.00500

[PD] Each wrong signal splits equally between the other states. All 27 signal-to-action policies are enumerated per scenario. No-information value is 1.75; EVPI is 2.45. Low/base/high are alternative likelihood and price judgments, not confidence intervals.

Bound the release decision

[PD] First audit whether the current file, affected dependency graph, capacity period and approval scope agree. Register the revalidation protocol, prices and strong comparators before testing. Preserve interim submission while a named authority resolves the gap; a completed workbook alone does not close it.

[PD] Coherent research-utility model, independent of an assumed programme benefit. No implemented control, corrected historical decision or customer effect is claimed.

REPRODUCIBILITY / READER PROTOCOL

Reproduce the decision basis, including the unresolved forks

[AI] VPEC is a proposed evidence contract. Its purpose is to reveal which decision is justified by the current record and which additional evidence could change that decision.

Layer	Reproduction requirement
Source lineage	Verify the source workbook hash, exact selected cells and claim locations. Separate recovered planning content from authentication of its registered version. Confirm that metadata is stored in the file; a reading library can supply default dates when properties are absent.
Reconciliation	Rebuild each portfolio scope, effort basis, value subtotal and metric population. Keep conflicting approval assertions unresolved unless authoritative evidence discriminates them.
Logical contribution	Reproduce the scheduling and critical-gate counterexamples. These are independent constructions; they establish insufficiency of a statistic, not an observed historical outcome.
Decision evaluation	Rebuild the version-change law, comparator policies, error costs, clustering, power and information value. Preserve adverse results against full review and interim submission.
Publication boundary	Raw values, internal identities and version coordinates remain in the private annex. Public hypothetical numbers are independent of confidential historical magnitudes.

[PD] An analyst should obtain the same conditional recommendation without choosing the favorable benefit version, changing prices to make the proposed policy win, or excluding unresolved original portfolios from the outcome denominator. A failed benchmark is a useful result if it changes the decision.

[AI] The private archive includes source manifests, exact calculations, the proposed evaluation law, review findings and rendered-page inspection records. Passing the publication process certifies the bounded argument and reproducibility; it does not establish that the historical plan was funded, executed or causally effective.

Protected source S1; case-specific source, numerical and methods reviews. Reproduction of source arithmetic is distinct from authentication of historical originals.

REFERENCES / 1

Primary sources and their role in this case

[AI] The literature supports specific analytical methods and governance mechanisms. It does not authenticate this project or transfer another study's effect to the portfolio.

NIST/SEMATECH. Individuals Control Charts, e-Handbook section 6.3.2.2. Official individuals-chart equations. Ordered draft versions and proposed targets do not become a stable operating phase merely because limits can be computed.

Laney, D. B. (2002). Improved Control Charts for Attributes. *Quality Engineering*, 14(4), 531–537. Use authentic counts and exposures for repeated rate subgroups. No missing count panel is imputed from rounded targets.

Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20, 37–46. Audit classification of approval scope, missingness and gate status with independent coders; agreement does not itself establish the correct authorization.

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. Selection of unusually troubled portfolios can create untreated reversion. Compare the frozen eligible frame and retain an explicit reliability sensitivity.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, 225, 200–230. For later observational adoption, define group-time effects with eligible not-yet or never-treated comparators; diagnose anticipation, spillover and pre-trend uncertainty.

Collier, D. (2011). Understanding Process Tracing. *PS: Political Science & Politics*, 44, 823–830. Separate reported authorship, accepted ownership, funding authority and execution. Registry-level role evidence does not prove realized benefit.

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 can change. Include its acquisition cost and preserve the option to retain the simpler policy.

Original author, institutional, official-method and publisher records. Access scope and case-specific relevance are recorded privately.

REFERENCES / 2

Primary sources and their role in this case

[AI] The literature supports specific analytical methods and governance mechanisms. It does not authenticate this project or transfer another study's effect to the portfolio.

The W. Edwards Deming Institute. The Funnel Experiment. Repeated intervention on routine variation can introduce instability. A sequence of plan revisions is not a performance-improvement time series.

Campbell, D. T. (1976). Assessing the Impact of Planned Social Change. A completion target can encourage filling fields without resolving authority. Preserve an independently adjudicated decision-validity outcome.

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 targeted indicator may lose its relationship to the objective. Readiness requires critical conditions, not only a high completeness score.

Ittner, C. D., & Larcker, D. F. (2003). Coming Up Short on Nonfinancial Performance Measurement. Harvard Business Review, 81(11), 88–95, 139. Validate the link from planning process measures to operating outcomes. A complete submission is an intermediate construct.

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. Local workarounds can leave recurring system constraints intact. A plan that repeatedly carries unresolved ownership needs a structural decision.

OECD / European Commission JRC (2008). Handbook on Constructing Composite Indicators: Methodology and User Guide. A compensatory weighted score can hide an indispensable failed gate. Preserve component definitions and test weighting sensitivity.

Phillips, L. D., and Bana e Costa, C. A. (2005). Transparent prioritisation, budgeting and resource allocation with multi-criteria decision analysis and decision conferencing. LSEOR 05.75. London School of Economics. ISBN 0-7530-1697-4. Application: Distinguish portfolio appraisal from the accountable allocation decision; make resource interactions and stakeholder assumptions explicit.

Original author, institutional, official-method and publisher records. Access scope and case-specific relevance are recorded privately.

REFERENCES / 3

Primary sources and their role in this case

[AI] The literature supports specific analytical methods and governance mechanisms. It does not authenticate this project or transfer another study's effect to the portfolio.

Johnson, A., Dempsey, K., Ross, R., Gupta, S., and Bailey, D. (2011; updated 2019). Guide for Security-Focused Configuration Management of Information Systems. NIST SP 800-128. doi:10.6028/NIST.SP.800-128. Application: Bind approvals to a recorded baseline, assess changes and verify the affected version before release.

U.S. Government Accountability Office (2020). Cost Estimating and Assessment Guide: Best Practices for Developing and Managing Program Costs. GAO-20-195G. Application: Use a common technical baseline, full cost scope, documented assumptions and dependence-aware uncertainty.

HM Treasury (2026). The Green Book: Central Government Guidance on Appraisal and Evaluation. Application: Retain unpriced effects explicitly and test switching values instead of treating missing valuation as zero.

Government Project Delivery (2025). Government Functional Standard GovS 002: Project Delivery. Version 2.1. Application: Specify approval authority, baseline changes, linked dependencies and accepted benefit ownership.

Hudgens, M. G., & Halloran, M. E. (2008). Toward Causal Inference with Interference. Journal of the American Statistical Association, 103, 832-842. Original manuscript. Reviewers, approval queues and shared dependencies can create spillovers. Randomize domains that contain those interactions and retain every original submission.

U.S. Government Accountability Office (2015). GAO Schedule Assessment Guide: Best Practices for Project Schedules, GAO-16-89G. Annual resource totals cannot substitute for a resource-loaded, logically connected schedule; test dependencies, timing and the effect of changed assumptions.

Original author, institutional, official-method and publisher records. Access scope and case-specific relevance are recorded privately.

FINAL REFEREE REPORT / INTERNAL REVIEW

Accept the planning diagnosis; keep the policy decision open

[AI] Author-led, agent-assisted review; not external peer review. Accepted scope: evidence reconciliation and a falsifiable planning-control design.

Major issue	Revision and disposition
Readiness and attribution	Separate submission, interim acceptance and execution lock. Bound personal contribution to registered analysis and coordination.
Version and population conflicts	Keep approval scope, staffing demand, row sets, missing-value types and financial bases explicit. Do not select a convenient headline.
Decision-changing inference	Show why equal annual capacity totals can be infeasible and equal completion scores can hide different critical blockers. Neither proof estimates historical impact.
Policy and economics	Retain snapshot, full-review and all-interim alternatives. The proposed targeted rule is not preferred at the base hypothetical prices; costs and ambiguity can change the ordering.
Evaluation adequacy	Retain inadequate smaller designs. Exact finite-null calculations resolve a simulation discrepancy; a high rejection rate in one run does not establish systematic inflation. Expanded recruitment remains a conditional candidate.

Partner critique: can a COO act on page one?

[AI] Yes. Accept an auditable interim submission; assign unresolved decisions to the roles that can close them; verify scope and version before execution lock. Do not deploy the more elaborate review policy solely because it outperforms an outdated snapshot comparator.

Remaining limits and minor issues

[AI] The technical limitations govern interpretation: original authentication, actual authority, missing measurement linkage and historical benefits remain unresolved. Prospective cost, reference and recruitment inputs need calibration. Public wording and exhibits preserve these distinctions without disclosing confidential numerical coordinates.

REVIEW DISPOSITION / PASS WITHIN STATED SCOPE

[AI] Publish the reconciled planning case and conditional evaluation. Historical approval, implementation, causal outcomes and realized savings remain unverified.

Disposition is supported by final evidence, methods, independent numerical, partner and actual-page reviews retained privately. No external endorsement is asserted.