

Preventive Action Design & Controls

Turn a forecast into an authorized move, an accepted owner and a testable result

An anonymized operating-design case study of a versioned action library: policy-linked instructions, cross-functional ownership, agent-assisted construction and an incomplete handoff into the operating application.

BLUF / BOTTOM LINE UP FRONT

Release executable action bundles.
Do not mistake catalogue growth for use.

RECALCULATED FINDING

Derived from registered totals: actions both locally included and labeled forward-skilled can represent zero to less than 6% of the final catalogue. Marginal counts cannot identify which actions qualify; the labels do not independently validate predictive skill.

CONTRIBUTION IN THE SOURCE RECORD

Requirements, source boundaries, acceptance criteria and integration constraints. Provisional registry-level hoop support; agent-assisted content and verification.

EVIDENCE STATUS

Registry-based design and reconciliation case. Original artifacts and field execution remain unauthenticated; no customer effect or realized savings is established.

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

EXECUTIVE ABSTRACT

Authorize the action where it can still change the outcome

[AI] The project's contribution is the design of a more specific, accountable action layer. Its unresolved operating constraint is whether an authorized instruction reaches an equipped owner before the prevention window closes.

CONTEXT | A forecast lacked a usable operating response

[HR] A delivery network's early-warning tool relied on generic, reactive instructions. The redesign added leading-signal explanations, role-specific moves, operating touchpoints, policy provenance, local variants and routing to functions able to resolve the underlying problem.

OBJECTIVE | Connect decision support to responsible execution

[PD] Authorize bounded action packages with a verified current instruction, available inputs, an accepted owner, a valid operating window and a traceable release. Forecast skill, instruction validity and customer benefit require separate evidence.

APPROACH | Specify the system and preserve its evidence boundaries

[HR] The lead set requirements and honesty gates, directed agent-assisted construction and constrained concurrent application writes. [AI] The register supports these roles provisionally; agent-session statements about independent verification do not prove the lead personally wrote or ran the checks.

FINDINGS | Catalogue coverage is not jointly qualified coverage

[RC] The catalogue tripled against its initial count. Conditional category bounds permit zero overlap between local inclusion and forward-skill labels, with an upper bound below 6% of the final library. [AI] Neither bound establishes deployed or beneficial actions; a persistence-based instruction can still be useful when its causal mechanism and authority are sound.

DECISION VALUE | Fund the missing release contract

[PD] The operations sponsor should authorize a limited evaluation after the platform and receiving-function owners sign every action's release record. Require zero unresolved mandatory authority, input, owner or release gates within the trial bundle; require separate financial and safety acceptance before scaling.

IMPLICATIONS | Keep the design contribution and test the operating effect

[AI] Preserve the policy and provenance discipline, but replace completeness scores with action-level joins. The proposed Action Release Contract links context, authority, delivery, ownership and evaluation; catalogue size alone cannot satisfy it.

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

EXHIBIT 1 / FROM CATALOGUE TO USE

Treat construction, local inclusion and field use as different states

[HR] The registered library grew while application integration remained partial. A local rendering check and an unchanged distribution package describe different artifacts; neither supplies operator-use telemetry.

CONSTRUCTION DOES NOT ESTABLISH OPERATOR USE

Historical register and analytical inference • distinguish evidence states

CATALOGUE Three times initial count	LOCAL APPLICATION Under half included	FIELD USE Not authenticated
--	--	--

A local render check and a distribution package are different artifacts.
The register reports that the distributed package remained pre-merge.

Recover exact versions and exposure records before describing deployment or adoption.

[RC] The final library contains three times the initial registered action count. [AI] Adding cross-functional actions changed the unit’s scope: growth measures content breadth, not equivalent frontline capability or net-new operating value.

Evidence state	What it establishes and what remains
[HR] Versioned library	A registered construction output with schema checks. Recover the original files before authenticating their contents and build history.
[HR] Local application	Partial inclusion and a reported rendered-page check. The record also says the distribution package remained pre-merge.
[AI] Operating use	No authenticated exposure, accepted work, timely execution or mature customer outcome panel. Absence of evidence is not evidence of zero benefit.

[PD] Give each action stable identity plus explicit states: draft, validated, authorized, released, acknowledged, executed and retired. Store the exact content hash and application version at each transition. An application-wide release label cannot prove every embedded instruction is authorized.

PLATFORM OWNER DECISION

[PD] Verify the actual distributed artifact and operator-visible action IDs before expanding the catalogue. Preserve legacy IDs with retirement records rather than retaining unsafe content solely to satisfy a superset check.

[HR/RC/AI/PD] S1:M01/M09/D05/D08 and narrative implementation. Public catalogue growth is relative. Application inclusion is not deployment or adoption.

EXHIBIT 2 / JOINT QUALIFICATION

Join action-level evidence before claiming preventive readiness

[RC] Category totals alone leave the overlap between local inclusion and forward-skill labels unidentified. [AI] The missing join is an operational decision requirement, not a rounding inconvenience.

MARGINAL TOTALS CANNOT IDENTIFY THE QUALIFIED INTERSECTION

Recalculated set bounds • common identities and snapshot assumed

LOCALLY INCLUDED ∩ RECORDED FORWARD-SKILL LABEL
Possible share of the final library: 0 to less than 6%

$$\max(0, A + F - N) \leq \text{intersection} \leq \min(A, F)$$

A = local set size; F = forward-labeled set size; N = library size

This is not validated skill, causal benefit or the share of all useful operating controls.

[RC] For a common library of size N, with A locally included actions and F forward-labeled actions, the intersection lies between $\max(0, A + F - N)$ and $\min(A, F)$. The recorded totals imply a lower bound of zero and an upper bound below 6% of the library. The calculation assumes matching identities and snapshots, which the original membership files must verify.

[AI] This is a bound on two recorded labels, not a count of release-ready instructions. Source-level forward labels may be inherited, conditional or stale; data permission, policy applicability, owner acceptance, time-window feasibility and actual distribution further restrict a particular release. Multiplying marginal percentages assumes an unsupported dependence structure.

Distinct question	Evidence that answers it
[PD] Can risk be anticipated?	As-of inputs, appropriate baseline, relevant horizon, validated trigger and calibration for the intended population.
[PD] Can the action help?	A defensible intervention mechanism, applicable authority, executable capacity and an outcome evaluation. Predictive lift alone cannot answer this.

SPONSOR DECISION
[PD] Recover the action-ID crosswalk and inspect every proposed trial action. A useful conventional control can remain eligible without incremental predictive lift; describe its mechanism honestly instead of upgrading its model label.

[RC/AI/PD] S1:M04/M09 plus private overlap calculations. Bounds are deterministic set bounds, not confidence intervals. Exact historical counts and category crosswalks remain private.

EXHIBIT 3 / AUTHORITY AND OWNERSHIP

Route accountability to the role that can change the mechanism

[HR] The redesign separated defect location from remedy ownership. [AI] Classification guides routing; it cannot grant authority, secure capacity or justify a performance consequence.

THE LOCATION OF A FAILURE DOES NOT IDENTIFY ITS OWNER

Historical design and proposed control • defaults allow documented exceptions

EXECUTION GAP
Enable + verify adherence

CAPABILITY GAP
Accept + route remedy

MEASUREMENT GAP
Repair + reclassify

Before activation: authority + tools + time + accepted receiving capacity
A role assignment in a document is not another function’s commitment.

Preserve exception reasons; do not equate the gap taxonomy with the response.

[HR] The library distinguished execution gaps, missing capability and measurement problems. It also documented exceptions to the default mapping between gap class and response. [AI] Therefore, “compliance means coach” and “capability means escalate” are defaults with conditions, not interchangeable taxonomies. An existing policy does not prove a local operator can implement its remedy.

Release check	Required operating evidence
[PD] Local control	The named role has current permission, tools, training and time to perform the specific step in this context.
[PD] Accepted escalation	The receiving function accepts the work class, response window, queue limit and exception owner. A role in a catalogue is not consent.
[PD] Fair accountability	Assess adherence only to authorized, feasible instructions; separate unavailable capability and source-data defects from avoidable noncompliance.

[AI] External-owner acceptance is unauthenticated. Routing work outward can move congestion to a shared queue and consume the prevention window. Naming a role does not complete the intervention.

RECEIVING-FUNCTION OWNER DECISION

[PD] Sign an action-level capacity and response agreement before activation. Record rejection, expiry and fallback so unresolved ownership remains visible to the sponsor.

[HR/AI/PD] S1:M02, governance and contribution records. No named external function is represented as having accepted or executed the proposed work.


EXHIBIT 4 / POLICY AND VERIFICATION

Separate a citation, an applicable instruction and a verified release

[HR] The design exposed missing sources and uncertain citations. [AI] A document identifier and a successful schema check cannot certify the recommendation’s meaning.

CITATION VALIDITY AND OPERATING AUTHORITY ARE DIFFERENT TESTS

Proposed release evidence • each layer answers its own question



Verify meaning independently of the authoring and schema-validation process.

[RC] The register’s cited-action and blank-identifier totals exceed the catalogue total if treated as disjoint categories. Its category summary also mixes verified, partial and flagged records. [AI] Preserve those as separate source assertions until a row-level crosswalk resolves membership; do not select the most favorable citation percentage.

[HR] Inaccessible instructions remained pending. Impact tags could coexist: a sourced defect rate could accompany an estimated preventable fraction. [AI] A sourced input does not make the product a sourced treatment effect. Overlapping tags cannot be added as action coverage.

Verification layer	Independent question
[PD] Syntax and provenance	Do fields parse, IDs resolve, cited versions exist and quoted passages match? Preserve the retrieved text and source timestamp.
[PD] Semantic applicability	Does the instruction authorize this actor, step, population and operating context now? Resolve exceptions with the accountable policy owner.
[PD] Delivered behavior	Does the released application display the approved version and block expired or inapplicable actions? Test actual serving paths.

POLICY OWNER DECISION

[PD] Review high-consequence instructions against current authoritative text. Keep inaccessible or disputed guidance pending; assign a dated recovery task rather than a fabricated instruction.

[HR/RC/AI/PD] S1:M03/M08/D03 and narrative evidence limitations. Original policy text is not authenticated in this case. A citation-status conflict constrains the claim, not just the footnote.


EXHIBIT 5 / TIMING AND CAPACITY

Put the entire response path inside the prevention window

[HR] The project moved instructions toward operating touchpoints. [AI] A nonweekly label does not establish a daily action or prove completion before a failure becomes irreversible.

A PREVENTIVE INSTRUCTION MUST FIT THE WHOLE RESPONSE PATH

Proposed time-budget test • measure signal, display and execution timestamps



A nonweekly label does not establish a timed daily routine or timely completion.

[RC] The cadence distribution includes untimed entries within its nonweekly count; nonweekly coverage therefore differs from daily timed coverage. [PD] Validate action-level clock windows, local time zones, input latency and shift handoffs.

[PD] Define usable lead time as irreversible-event time minus signal availability. Authorization delay, routing delay, queue wait and execution time must fit within that interval, with a prespecified buffer. If the window closes, the interface should change the action’s state and show an authorized corrective fallback rather than keep displaying a preventive promise.

Mechanism	Failure mode the release must expose
[AI] Shared specialist owner	Demand spikes can make a sound instruction infeasible. Queue saturation can transmit spillovers across sites.
[AI] Delayed source refresh	A nominally early signal may arrive after the local control opportunity. Validate event time, ingestion time and display time separately.
[PD] Capacity gate	Measure arrival and service distributions by work class. Stop expansion when timely-completion or safety guardrails fail.

OPERATIONS OWNER DECISION

[PD] Instrument timestamps and accepted work before scaling. An initial trial bundle should contain actions with explicit response capacity, observable completion and a mature outcome definition.

[HR/RC/AI/PD] S1: narrative governance, implementation-governance cadence distribution and source exceptions. No historical latency distribution or queue-service performance is available.

EXHIBIT 6 / IMPACT REPRESENTATION

Keep scenario arithmetic separate from causal benefit and calibrated uncertainty

[HR] The application used modeled magnitudes and bands. [AI] Regression links and assumed achievable fractions do not identify the effect of executing an action.

MODELED REDUCTION IS AN ASSUMPTION CHAIN, NOT A MEASURED EFFECT
Historical model structure and analytical inference • protect interpretation

DEFECT MEASURE
Observed or forecast

ADDRESSABLE SHARE
Assumed

PASS-THROUGH
Inherited relationship

RECONCILIATION
Snapshot-specific

A coherent arithmetic total can still have an unidentified treatment effect.
Do not sum overlapping actions or label uncalibrated scenarios as quantiles.

Evaluate the assigned bundle once; value the mature outcome on its own cost basis.

[HR] The chain combines a defect measure, assumed addressable fraction, inherited coefficient and snapshot reconciliation factor. [AI] Reconciling a headline cannot validate preventability, identify an intervention effect or eliminate overlapping action benefits.

[RC] One signed example places P10 numerically above P90. [AI] Applying a minus sign to positive reduction quantiles would reverse their order for the signed change. Until the random variable and calibration are documented, publish a directional scenario range without probabilistic coverage.

Model choice	Release rule
[PD] Snapshot identity	Freeze the bucket, headline and reconciliation factor together. Run alternate snapshots separately; never combine their denominators.
[PD] Portfolio overlap	Estimate the assigned bundle’s mature effect once. Attribute components only with additional design; do not sum overlapping per-action deltas.
[PD] Cash conversion	Apply avoidable cost and non-overlap assumptions to the same mature eligible population, then subtract incremental costs.

FINANCE OWNER DECISION
[PD] Treat historical magnitudes as scenario inputs. Require a prespecified economic hurdle and an uncertainty bound on the evaluated bundle before claiming net operating value.

[HR/RC/AI/PD] S1:M06/M10/M11/D06 and private signed-band check. No historical probability calibration, treatment effect or realized cash benefit is established.


EXHIBIT 7 / ACTION RELEASE CONTRACT

Make every mandatory release condition explicit and owned

[PD] The Action Release Contract is a proposed control for turning a versioned instruction into a testable operating intervention. Its gates are conjunctions; strong documentation cannot compensate for absent authority or unavailable delivery.

THE ACTION RELEASE CONTRACT JOINS FIVE MANDATORY CONDITIONS

Proposed ARC framework • no weighted average replaces a missing gate



Same action identity. Same approved version. Same population and time.
Reopen the affected gate when any material contract field changes.

Proposed release discipline; not an observed rollout, certification or causal result.

Contract field	Evidence and accountable role
[PD] Context and mechanism	Operations: intended unit, trigger, causal rationale, valid window and corrective fallback.
[PD] Authority and evidence	Policy owner: current applicable instruction, scope, source status, permissions and exception disposition.
[PD] Delivered version	Platform: exact content hash, input availability, serving-path test, feature state and reversible release.
[PD] Accepted ownership	Receiving function: capacity, response commitment, acknowledgement and escalation on expiry.
[PD] Outcome evaluation	Analytics and finance: assigned population, mature endpoint, estimator, economic hurdle and separately powered harm guards.

NEXT INVESTMENT

[PD] Recover action IDs, authorize a limited bundle, verify distribution and secure owner acceptance. Evaluate that package before expanding its scope.

[PD/AI] Proposed ARC framework, grounded in S1 decision and implementation constraints. It is neither an observed rollout nor a certified control standard. Falsifiable propositions appear in the technical layer.


TECHNICAL FOUNDATION • THEORY OF CHANGE

Availability changes an operating pathway; the customer effect remains untested

[PD] The Action Release Contract (ARC) separates context, authority, local delivery, accepted ownership and outcome testing. Its proposed causal graph distinguishes a usable recommendation from a delivered intervention.

ASSIGNMENT AND EXECUTION ARE DIFFERENT CAUSAL EXPOSURES

Proposed DAG • arrows are hypotheses; historical causal effects are not identified



Shared queues link domains unless randomization follows the actual interference boundary.

ITT follows all paths from Z to Y; conditioning on execution can introduce selection bias.

[PD] Z assigns availability of the frozen authorized package plus orientation. A is an accessible, applicable recommendation; T is execution; M is the operating mechanism; Y is the prespecified customer failure within 28 days of the original promise. C represents measured capacity and case mix; U represents unmeasured severity and concurrent work.

[AI] The registered design concerns A and the requirements for A→T; randomized Z belongs to the proposed evaluation. No historical estimate identifies T→M→Y. Authority, timing and owner acceptance can interrupt this path even when a card exists. The separate Z→Y path abbreviates other operating effects through monitoring or routing.

[PD] Record outcome ascertainment separately: Y produces recorded outcome R through a measurement process Q. Missing feeds, delayed reporting and taxonomy errors can alter R without changing Y. Residual differences can therefore arise from unmeasured operating paths, measurement error or chance; a residual does not identify its own explanation.

[AI] NIST's attribute-based authorization model provides a design analogy for context-specific permission. It does not certify this operating library. Recommendation generation, decision authority and execution remain distinct functions, consistent with Parasuraman, Sheridan and Wickens.

Proposed DAG; registered requirements are historical evidence of design, not estimates of the causal arrows. NIST SP800-162; Parasuraman et al. (2000).

TECHNICAL FOUNDATION · ESTIMANDS AND DESIGN

Randomize access to the package and retain every eligible unit in the outcome

[PD] The primary estimand is the effect of offering the authorized package, not the effect among people who choose to execute an action.

[PD] Enroll eight weekly cohorts of upstream-eligible customer units in independent service domains. Freeze eligibility and original promise time before treatment can affect selection. Assign domains to the package plus standard orientation or usual approved practice, with the same forecasting feed and baseline staffing. Keep holds, cancellations and failed deliveries in their assigned population; do not restrict analysis to completed work or opened cards.

Quantity	Population and interpretation
Availability ITT	Usual-practice minus package-offer 28-day failure risk; positive means benefit. Keep all original eligible units, with the prespecified domain weights.
Execution effect	Effect of actually delivering a specified action package. Not identified by comparing voluntary users with nonusers.
Mechanism evidence	Timely execution, authority compliance, owner acceptance and capacity use. These explain fidelity; they do not substitute for customer outcomes.

[PD] Read each unit through 28 days after its original promise; the final readout follows the last enrolled promise by that full window. This is a defined 28-day endpoint, not an assertion that all eventual events have matured. Blind outcome coding to assignment, freeze the defect taxonomy, reconcile reporting delays and provide missing-outcome bounds. Estimate the trial contrast using the same population weights as its economic model.

[AI] Orientation and routing can affect outcomes without the recorded action, so the exclusion restriction for an instrumental-variable execution effect is doubtful. Dividing ITT by uptake is not justified. Angrist, Imbens and Rubin specify the additional assumptions; clustered encouragement designs also require cluster-aware inference.

[PD] Randomize at the shared staffing and support-capacity domain. If a queue serves both arms, enlarge the cluster or redesign. Keep concurrent controls through outcome readout; early crossover would change the comparison. Prespecify business-line and node-type strata, baseline trend diagnostics, contamination measures and harm guardrails. Baseline similarity does not prove absence of bias. The numerical power design is conditional planning; correlated action rows cannot replace independent domains.

Proposed parallel cluster availability trial. Angrist et al. (1996); Frangakis et al. (2002); Moore et al. (2015). No historical treatment or customer effect is asserted.

Separate evidence dimensions before deciding that an action is ready

[AI] Catalogue coverage, policy citation, predictive labels, local rendering and owner acceptance measure different constructs. Their marginal percentages cannot be multiplied into a readiness score.

Construct	Validation contract
Authority	Retrieve the applicable rule, version, geography, process and delegated decision right. An identifier or inaccessible document is insufficient.
Local availability	Join immutable action IDs to the delivered build and operator view. Distinguish a local file, a packaged release and observed use.
Accountability	Record acceptance by a role with control, capacity and an escalation clock. A named function is not an accepted commitment.
Impact	Keep sourced defect magnitude, assumed preventability, modeled effects and measured mature outcomes in separate fields.

[HR/AI] The source describes overlapping source and estimate tags: an authentic input can coexist with an assumed gain. It also contains policy-status totals that do not form a clean exclusive partition. Until row-level reconciliation establishes scope and precedence, neither set of labels supports a complete distribution of verified actions.

[RC/AI] Conditional on matching library versions, the intersection of locally represented entries and entries carrying a forward label ranges from zero to under 6% of the catalogue. This is a set-overlap bound, not measured readiness. The label itself does not prove forward predictive skill, and the bound does not cap the usefulness of persistence-based or deterministic actions.

[PD] Double-code a proposed random core of 200 action-context decisions, stratified by business line, node type and source status. This samples contexts, not 200 unique current cards. An independent-observation calculation gives a worst-case raw-agreement margin of about ±6.9 percentage points at 95%; shared cards and policies require a design-effect adjustment. Pilot category prevalence and disagreement patterns must determine kappa precision by simulation. Report confusion matrices and ratings before adjudication against a frozen codebook. Supplement rare exceptions separately; population estimates require appropriate sampling weights.

[PD] Test invariance with equivalent operating vignettes across groups: identical labels must imply the same authority and failure construct. Version joins, denominators and missingness remain explicit. W3C provenance can document derivation and responsibility; it does not establish that a claim is true.

Registered metadata and private set reconciliation; proposed measurement protocol. No historical row-level readiness join is supplied. W3C PROV-DM (2013).

TECHNICAL FOUNDATION • PREVENTIVE MECHANISMS

An action is preventive only while authority, capacity and timing still permit it

[AI] A nonweekly cadence is not a verified preventive window. The source's nonweekly category includes entries without an assigned execution time.

[PD] Define the usable window from the first valid as-of signal to the last reversible operating decision. Require data availability, queue waiting, owner response and execution latency to fit inside that window. Record the actual expiry event rather than borrowing a meeting cadence. A recommendation that arrives before an outcome is reported can still arrive after prevention is impossible.

[HR/AI] The registered taxonomy distinguishes compliance, capability and measurement gaps, with documented exceptions to its usual coaching and shielding routes. A clean rule check therefore means no unapproved mismatch under that taxonomy; it does not prove a universal one-to-one correspondence. An unavailable policy document also does not prove that no applicable requirement exists.

[PD] Coach only where the actor has the relevant obligation, resources, control and opportunity. Route a capability constraint to the function able to change it, preserving the frontline protection rule and a named acceptance deadline. Exceptions need a rationale, approving authority, scope, expiry and review record. Measure orphaned escalations and displaced support work so transferring ownership cannot conceal an unresolved constraint.

[AI] The source's impact chain combines a defect magnitude, assumed achievable fraction, statistical passthrough and reconciliation factor. A regression association cannot turn that product into an action effect. Shared defect pools make per-card effects nonadditive. Shmueli's distinction between prediction and explanation is directly relevant; low predictive fit also does not establish that an intervention cannot work.

[PD] Use positive avoided-failure units or explicitly signed changes throughout. Percentile labels require a defined distribution and correct ordering; otherwise show the range as an assumption sensitivity. Translate outcomes to value only after matching population, maturity and cost definition, removing overlap and subtracting incremental delivery costs. A gross unit-cost citation does not establish avoidable expenditure.

[AI] The new design implication is a coupled clock: the action's expiry and the receiving owner's acceptance deadline must be jointly feasible. A longer catalogue can increase expired or orphaned work if support capacity stays fixed. That mechanism is testable, not an observed project result.

Registered cadence, taxonomy and modeled-impact descriptions; proposed window and capacity tests. Shmueli (2010); Tucker and Edmondson (2003).

TECHNICAL FOUNDATION • VARIATION AND SELECTION

Artifact growth does not establish process improvement or treatment response

[AI] Versioned catalogue counts, metadata profiles and inherited operating anchors are not a stable, repeated outcome panel for this intervention.

[PD] For a comparable ordered series, the disclosed individuals-chart calculation is $MR(t) = |x(t) - x(t-1)|$, with tentative limits $\text{mean}(x) \pm 2.66 \times \text{mean}(MR)$ and moving-range upper limit $3.268 \times \text{mean}(MR)$. Short-series limits remain uncertain; a constructed two-point pair cannot establish a stable process. The private audit distinguishes mechanical calculations from eligible operating evidence.

[HR/AI] Changes in the registered catalogue reflect additions, consolidation and new fields. Different categorical labels are not consecutive time points. Inherited outcomes have changing scopes and definitions, and the source explicitly flags a suspected measurement artifact. These records cannot identify action-library improvement. No commensurate event-count/exposure panel supports a Laney-adjusted rate chart for this project.

A worst-performing cohort can improve without treatment

[PD] No fixed worst-N action-treatment cohort with linked follow-up is supplied. A synthetic regression-to-the-mean illustration sets a selected baseline index to 100 and its population mean to 80. Under a stationary equal-variance Gaussian repeat-measurement model, $E[\text{next} | \text{baseline}] = 80 + \rho \times 20$. Reliability values 0.25, 0.50 and 0.75 imply expected next indices 85, 90 and 95 without any intervention. These are assumed scenarios, not historical adjustments.

[PD] A real evaluation must freeze the selection date, ranking denominator, eligibility and all selected units. Compare focus and nonfocus units over the same calendar and maturation windows, using pretreatment histories to estimate persistence and case mix. Do not interpret within-limit movement as improvement; do not treat a chart crossing as treatment attribution.

[PD] Prespecify business line, node type, preventive-window length and support dependency as moderators. A locally controlled routine may respond to coaching; an externally constrained process may require a capacity remedy first. Compare interaction estimates with uncertainty, not significance in one subgroup versus nonsignificance in another. Sparse strata warrant partial pooling or a stated evidence gap.

No historical effect estimated. NIST individuals-chart method; synthetic mean-reversion illustration derived from the explicitly stated repeat-measurement model.

Credit the specification and acceptance logic without claiming unsupported execution

[HR/AI] The register supports a project leader's role in direction, requirements and review at a provisional level. It does not authenticate personal authorship of every entry or operational adoption.

Contribution claim	Process-tracing grade and next evidence
Defined preventive requirements and acceptance rules	Provisional hoop: the register attributes direct specification decisions. Authenticate the original request, version and responding artifact.
Directed evidence harvest and challenged unsupported claims	Provisional hoop: described review and correction activity is consistent with that role. Seek independently attributable review records.
Personally executed each script or authored every card	Not established. Agent/session first-person language does not authenticate a human actor; automated drafting is separately described.
Delivered field adoption or customer improvement	Not established. Local integration and verification descriptions do not supply accepted owners, execution receipts or mature comparisons.

[AI] A hoop test asks whether necessary evidence for a claim is present; passing it does not exclude competing explanations. A smoking-gun test would need a uniquely attributable, contemporaneous artifact connecting the leader's decision to an implemented change. Doubly decisive evidence would also eliminate credible alternatives. This record does not reach those stronger grades.

[PD] Preserve the human specification, agent-produced draft, reviewer changes, script version and build artifact as distinct entities. Retain timestamps, authorship provenance and handoff receipts. Link a directed requirement to the exact accepted diff, then link that build to operating use. Repeated summaries of the same underlying record do not provide independent corroboration.

[HR/AI] Reported verifier success concerns its declared check set. It cannot authenticate an inaccessible policy, establish receiving-function agreement, or prove that operators used the result. An approved exception may pass a rule while still needing independent semantic review. No adoption trace was recovered; this does not establish whether adoption occurred elsewhere.

[AI] The defensible contribution is the architecture of an accountable preventive catalogue and its quality requirements. Parent forecast performance, a sibling application's deployment and another site's results require separate attribution. Keeping these boundaries visible strengthens the operating case rather than reducing the importance of the design work.

Registered project narrative, decisions and evidence metadata; provisional process-tracing assessment. Original-source recovery limits are consolidated in the limitations section.

TECHNICAL FOUNDATION • TESTABLE CONTRIBUTION

ARC makes five release conditions testable rather than treating completeness as success

[PD] ARC is a proposed operating framework: context, authority, local delivery, accepted owner and outcome test. It has not been validated as a general theory or a proven intervention.

P1 • Contextual authority improves recommendation validity

[PD] Against a citation-only comparator, explicit actor, geography, process and delegation checks reduce unauthorized recommendations on a frozen, independently adjudicated vignette set. Retain all eligible requests and track valid-action coverage and review time. A precise increase in invalid recommendations, or failure of the coverage guardrail, contradicts the proposed advantage.

P2 • Accepted ownership reduces unresolved handoffs

[PD] Requiring capacity acceptance and an expiry clock reduces orphaned or expired escalations relative to assigning a function name. Compare the same workload and support budget; count transfers, rejections and unresolved work. A sufficiently precise increase in unresolved workload contradicts the mechanism. Mere receipt of an escalation does not count as remedy completion.

P3 • Preventive timing predicts delivery before lockout

[PD] Recommendations with a feasible signal-to-action window achieve more independently verified pre-expiry execution than equally valid recommendations lacking a feasible window, conditional on prespecified capacity and case mix. This process association tests a timing hypothesis; causal attribution needs randomized timing or additional identification. Complete timestamps and a stable lockout definition are required.

P4 • Availability creates customer value at acceptable cost

[PD] The authorized package improves the randomized 28-day customer outcome enough to clear the same-population economic threshold without unacceptable harm. A lower confidence bound above break-even supports scaling; an interval spanning the threshold is inconclusive. An upper bound below break-even contradicts the economic proposition under its stated cost assumptions. The bundle comparison does not identify each card's contribution.

[AI] The contribution combines contextual authorization with implementation fidelity and boundary-spanning learning. NIST informs permission structure; Moore and Carroll distinguish delivery from outcome; Tucker and Edmondson explain why immediate workarounds may fail to remove recurring constraints. ARC extends those ideas into a proposed release contract, not a claim that their findings have been replicated here.

Four prospective falsifiable propositions. No historical validation is claimed. NIST SP800-162; Moore et al. (2015); Carroll et al. (2007); Tucker and Edmondson (2003).

TECHNICAL FOUNDATION • LITERATURE AND DECISION LOGIC

The ranking rule must measure decision value rather than reward the easiest label

[AI] The case connects established measurement problems to a specific design risk: reliable evidence about a baseline can be mistaken for reliable evidence about an action's benefit.

Source priority is an explicit value judgment

[HR/AI] The registered ranking places the source tier before the impact magnitude; it is a lexicographic ordering, not a weighted score or merely a tie-break. Under its sourced-wins rule, a mixed-tag card can gain priority from a documented baseline while its preventable effect remains assumed. Preserve field-level lineage and audit sensitivity to the precedence rule. This exposes a construct mismatch; it does not prove the selected actions were inferior.

Completeness can become a distorted target

[AI] Campbell's indicator argument predicts pressure to optimize a measure when it controls decisions. Here, rewarding citation presence or filled fields could favor apparent completeness over applicable authority. That is a prospective incentive risk, not evidence of misconduct. ARC therefore combines immutable source records with independent semantic checks and valid-opportunity coverage, including rejected and unresolved cases.

More responses are not automatically better control

[AI] Deming's funnel illustrates how adjustments to unstudied variation can worsen a process. ARC does not prescribe inaction: a binding safety rule or deterministic capacity constraint can justify action without a statistical signal. It requires the response to match the mechanism, distinguishing noise-driven recalibration from authorized prevention. The available artifact versions cannot test whether historical tampering occurred.

Measurement quality and economic performance remain separate

[AI] Ittner, Larcker and Randall found that measurement-system satisfaction and economic performance need not track the same management practices. The case likewise cannot infer value from a more elaborate catalogue. OECD/JRC composite guidance makes weighting, aggregation and sensitivity explicit; ARC applies that discipline to source precedence and to noncompensatory release gates. A strong score cannot compensate for absent authority.

[PD] The proposed extension is to test both ranking stability and realized decision value on a fixed eligible population. Change the source-priority rule transparently, retain the same capacity and authority constraints, and compare the chosen action set before any outcome claim. The present evidence diagnoses a design vulnerability; it neither confirms nor contradicts a field-effect theorem.

Registered ranking rule; proposed application of Campbell (1976), Deming's funnel, Ittner et al. (2003), and OECD/JRC (2008). No historical incentive or financial effect is inferred.

TECHNICAL FOUNDATION • LIMITATIONS AND REPRODUCIBILITY

The next evidence must close release and outcome gaps, not enlarge the catalogue

[AI] The evidence supports a registered design case. It does not establish an authenticated deployed library, population-wide policy validity, measured execution or realized customer benefit.

[HR/AI] The selected evidence is a secondary project register; the underlying source artifacts were not recovered for this case. Registered versions describe different build states and scopes. Cross-domain sources establish candidate mechanisms, not local authority or transportable effects. Status conflicts and overlapping tags remain explicit rather than being silently forced into a clean partition.

[PD] Reproduction requires a versioned row-level manifest linking action ID, source rule, business line, node type, model linkage, delivered build, owner acceptance, execution receipt and outcome cohort. Freeze the codebook and count each eligible record once in its declared denominator. Archive both the selected source interpretation and alternatives; rerun set bounds and conclusions under each compatible interpretation.

[PD] Prospective outcome collection must separate event occurrence from reporting and preserve all upstream eligibility. Report arm-specific missingness and plausible worst-case bounds. Test whether taxonomy, capture and maturity differ across groups. If the action changes reporting behavior, independently audited outcomes are needed to separate measurement effects from customer effects.

[PD] The planning model's variance, uptake, preventability, available domains and unit costs are assumptions until a pilot supplies evidence. Independent domains may be scarce when support queues overlap. If recruitment cannot meet the required power, redesign the horizon, endpoint or delivery model and recalculate; repeated cards or weeks do not create independent capacity domains.

[PD] Validate outcome-to-value conversion with matching units, severity mix, avoidable-cost share, benefit overlap and incremental package cost. Report capacity value separately from cash. Value of information concerns whether further evidence changes a decision and is not added to the same operating benefit. Historical realized value remains unestimated.

[AI] The appropriate next authorization is a bounded release-validation exercise followed by the prespecified randomized availability test. Existing permission and safety requirements apply to both arms. A favorable economic scenario cannot override missing authority, and a successful technical check cannot authorize an unsupported operating action.

Consolidated limitations and proposed reproducibility requirements. Raw source values, version alternatives and calculations remain in the private annex; public evidence is anonymous and normalized.

PROSPECTIVE MODEL / ASSIGNMENT AND COVARIANCE

Size independent operating domains, not the number of action cards

[PD] Every numerical input in this planning model is hypothetical. The exercise tests whether a proposed evaluation could resolve an economic decision; it supplies no achieved project power.

Population, outcome and treatment contract

[PD] Randomize independent service domains to a frozen authorized library plus standardized orientation or usual approved practice. Shared operators, escalation capacity and queues belong within one domain. Use the same forecasts and baseline staffing. Retain all upstream-eligible units across eight weekly enrollment cohorts; record the union of loss, damage, late delivery or nonfulfillment, counting each unit at most once.

[PD] Observe each unit through day 28 after its original promised event, including held, cancelled and unfulfilled units. The final cohort receives the full follow-up. Missing ascertainment is not a negative outcome. The endpoint is finite 28-day risk, not an assertion of ultimate maturity.

Planning input	Hypothetical value and meaning
Reference and effect	Untreated reference risk index=100; usual-minus-offer ITT=12 index points. One point is 1% of the frozen reference risk, not an absolute percentage point.
Domain covariance	Weekly SD=18; persistent correlation component=0.30; residual AR(1) coefficient=0.40. Independent domains; equal exposure by domain and cohort.
Estimator	Difference between arm averages of domain means over eight cohorts. Constant treatment effect in the planning alternative; equal reference and severity mix.

The covariance does not disappear with more observations

[RC/PD] For weeks t and s , $\Sigma(t,s)=18^2[0.30+0.70\times0.40^{|t-s|}]$. Averaging its entries gives domain-mean variance 155.48016, so $SE=\sqrt{2\times\text{variance}/n}$, with n domains per arm. The Gaussian rate-index distribution is an approximation, not a literal bounded event process.

[PD] Before recruitment approval, a blinded pilot must establish exposure weights, covariance, interference boundaries and the feasible number of independent domains. Re-power the actual estimated-variance analysis, including studentized or randomization-based sensitivity; the known-covariance calculation is not a field guarantee.

[PD/RC] Hypothetical Gaussian aggregate model. `private/model.py` and independently recomputed outputs. The endpoint, estimator, index and economic reference pool are declared together.

PROSPECTIVE MODEL / DETECTABILITY AND THE DECISION

A test that detects a nonzero effect may still fail the investment decision

[PD] Require both 80% power to detect break-even against zero and 80% probability of clearing break-even with the lower 95% confidence bound at the planning alternative.

[RC/PD] Under the base cost assumptions, break-even is 5.88235 reference-risk points. With true ITT fixed at 12 points, the economic decision statistic is the lower endpoint of a two-sided 95% interval. A p-value against zero answers a different question.

Domains / arm	80% MDE	Economic power	Joint gates
36	8.23	54.84%	Fail
72	5.82	83.74%	Conditional pass
84	5.39	88.87%	Conditional pass

[RC/PD] The smallest integer base plan satisfying both requirements is 71 domains per arm (142 total). A balanced 72-per-arm plan is a convenient candidate; 84 per arm adds conditional statistical margin. None establishes feasible recruitment.

A plausible covariance change can remove the margin

Stress at 84 / arm	Result	Required minimum / arm
Higher noise (SD=22)	MDE 6.59; economic power 73.94%	106
Higher persistent correlation (0.45)	MDE 5.98; economic power 81.71%	87

[PD] If independent domains are unavailable or the pilot indicates higher noise, redesign before starting: improve outcome precision, alter the supported service pool, extend genuinely informative follow-up or secure additional independent capacity domains. Recalculate costs and power together. Adding correlated weeks or cards does not automatically resolve the deficit.

NO AUTOMATIC GO DECISION

[PD] Size safety, workload, unequal-treatment and invalid-instruction guardrails separately, with prespecified margins and multiplicity. A financial power calculation does not demonstrate acceptable harm risk.

[RC/PD] Conditional known-covariance planning; two-sided $\alpha=0.05$. Minimum detectable effect and break-even use the same index. Final estimated-variance trial sizing remains a pilot-dependent gate.

PROSPECTIVE MODEL / SAME-POPULATION ECONOMICS

Value the assigned package once and keep the loss pool fixed

[PD] The financial model uses the mature availability-package ITT from the same eligible population. No additional uptake multiplier is applied to an effect that already includes nonexecution.

[PD] Set the fixed gross-loss pool to 100 for the intended service volume, eight enrollment cohorts and complete 28-day follow-up. Assume a 0.60 avoidable-cost share, a 0.85 exclusive financial-credit share and incremental operating/change cost of 3. These assumptions are illustrative; the register supplies no empirical values for this trial.

ECONOMIC IDENTITY

[RC/PD] Net value = ITT points × avoidable share × exclusive share – cost. Base net = 12 × 0.60 × 0.85 – 3 = 3.12. Break-even = 3 / (0.60 × 0.85) = 5.88235 points.

ITT points	Avoidable 0.40	Avoidable 0.60	Avoidable 0.80
6	-0.96	0.06	1.08
9	0.06	1.59	3.12
12	1.08	3.12	5.16

[PD] The grid varies causal-effect and cost-release assumptions; exclusivity and incremental cost stay fixed. These scenarios are not confidence intervals and need not be independent in practice. A larger physical effect may require more costly execution or reach lower-value defects, so correlated sensitivities should be elicited before investment.

[PD] Avoidability values expenditure or capacity that can be released; it does not discount physical efficacy twice. Exclusivity removes monetary items already credited elsewhere; use 1 when the pool is already exclusive. Stable expected loss per failed unit is assumed. If the intervention changes severity mix, use a cost-weighted endpoint and re-power it.

[PD] Allocate implementation, orientation, execution, monitoring and receiving-function costs to the same service pool. Trial recruitment does not redefine that volume. Research-only evaluation cost is separate and must enter the investment budget once. Capacity, avoided expenditure and verified cash release remain distinct claims.

[PD/RC] Synthetic same-pool sensitivity. Historical gross unit cost is retained privately and is not substituted for avoidable marginal cost. Bundle ITT avoids summing overlapping per-action modeled deltas.

PROSPECTIVE MODEL / SEPARATE INFORMATION VALUE

Buy evidence only when it can change a worthwhile decision

[PD] This separate decision model asks whether further evidence is worth buying before committing to a trial. Its utilities are illustrative option values, not measured savings or additions to the operating benefit.

Action / state	High potential	Needs preparation	Poor transfer
Fund controlled trial	20	-8	-22
Prepare and reassess	5	10	-6
Defer	0	2	4

[PD] State probabilities are 0.40, 0.35 and 0.25. Payoffs include downstream action costs. Without new information, prepare-and-reassess has the largest expected utility, 4. A signal identifies the true state with probability q and distributes errors equally over the other states. Priors, accuracy and payoffs must be elicited before operational use.

Signal q	EVSI	Evidence cost	Net EVSI
0.45	0.00	0.50	-0.50
0.65	2.88	1.25	1.63
0.85	6.09	2.50	3.59

[RC/PD] EVSI equals expected best posterior utility minus best current utility. EVPI is 8.50; every scenario satisfies $0 \leq \text{EVSI} \leq \text{EVPI}$. The weakest signal leaves preparation optimal for every possible result, so it has zero gross decision value despite returning information.

[AI] The source records validation holds, staged content and refused integration paths. It does not identify their realized monetary value. The model illustrates a way to price the next evidence decision without assuming a positive historical customer effect or relabeling an avoided hypothetical loss as observed savings.

DECISION BOUNDARY

[PD] Compare net information value with the evidence budget. Mandatory authority and safety gates remain prerequisites. Never add these option-value utilities to the trial’s operating benefit when both concern the same downstream decision.

[PD/RC/AI] Finite-state Bayesian decision illustration. State priors, likelihoods, action utilities and acquisition costs are hypothetical; exact arithmetic and posterior decisions are retained privately.

Reproduce the conclusion while preserving the publication boundary

[AI] The case is an operating-design and evidence-reconciliation study. Its public findings are bounded by what the protected register records and what the arithmetic can establish.

Claim class	Meaning in this case
[HR] Historical record	A registered requirement, decision, artifact state or reported check. Original library files and execution records are not authenticated by this label.
[RC] Recalculated	Set bounds, count reconciliation, signed-order checks or explicitly hypothetical planning-model arithmetic. Calculation does not convert an input into a causal result.
[AI] Analytical inference	A bounded interpretation of authority, timing, ownership, evidence quality or release status. Alternative explanations and unresolved source conflicts remain visible.
[PD] Proposed design	The Action Release Contract, evaluation protocol, synthetic economics and information-value model. These are recommendations, not historical implementations.

The reproduction order

[PD] Verify the frozen register-workbook hash and cell-linked extract. Recompute the set bounds and alternate source interpretations. Separately rerun the hypothetical trial and decision models using their declared assumptions. Recover the underlying historical action membership and versions before replacing bounds with a claimed intersection.

[PD] The private package retains source identities, raw inputs, locators, reconciliation alternatives, code, calculation outputs and review records. The public case reports protected quantities and proposed normalized scenarios. Do not distribute the private files with the public PDF.

[AI] The original contribution is the recognition that coverage statistics do not identify jointly qualified actions, and that the transition from description to authorized execution requires a version-specific contract. The bounds remain informative without asserting that any instruction was ineffective.

CURRENT EVIDENCE BOUNDARY

[AI] The registry-based design can be assessed and its arithmetic reproduced. Original artifact contents, field adoption, individual execution effects and realized financial value remain unestablished.

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

RESEARCH REFERENCES / METHODS AND INTERPRETATION

Distinguish an accurate prediction from an authorized intervention

Primary scholarly and official sources inform the method. They do not independently corroborate this project or validate its proposed operating framework.

Hu, V. C., Ferraiolo, D., Kuhn, R., et al. (2014; updated 2019). Guide to Attribute Based Access Control Definition and Considerations. NIST SP 800-162. Official guidance. Actor, action, object and context inform authorization. ARC uses this as a design analogy for operating permission; it is not an access-control implementation or compliance certification.

NIST (2023). Artificial Intelligence Risk Management Framework, AI 100-1. Official framework. Document roles, responsibilities, intended context and human oversight. Those requirements motivate explicit owner acceptance and a release scope; the framework supplies no evidence that this project achieved them.

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A Model for Types and Levels of Human Interaction with Automation. IEEE Transactions on Systems, Man, and Cybernetics A, 30, 286–297. Primary publication metadata and abstract. Automation of analysis, decision and action are distinct design choices.

Shmueli, G. (2010). To Explain or to Predict? Statistical Science, 25, 289–310. Author manuscript. Predictive and explanatory objectives differ. Forecasting an outcome does not identify how changing a process will affect that outcome.

Kleinberg, J., Ludwig, J., Mullainathan, S., & Obermeyer, Z. (2015). Prediction Policy Problems. American Economic Review, 105(5), 491–495. Primary publication. Some decisions benefit from improved prediction when the action-payoff relationship is understood. This prevents the opposite error of demanding a new causal model for every useful conventional control.

W3C (2013). PROV-DM: The PROV Data Model. Official recommendation. Preserve entities, activities, derivations and responsibility across versions. A documented derivation supports auditability; it does not prove the content is correct or an attributed operating action occurred.

Thurgood, S., with Frame, J., Lenton, A., et al. (2018). Alerting on SLOs. The Site Reliability Workbook. Original practitioner guidance. Alert quality includes precision, recall, detection and reset time. Its actionability principles inform the response-window test; software error budgets are not substituted for delivery outcomes.

Source scope is stated with each reference. Historical project evidence is the protected register S1; methodological sources establish interpretation and proposed designs, not project results.

RESEARCH REFERENCES / METHODS AND INTERPRETATION

Evaluate delivery, measurement and customer outcomes separately

Primary scholarly and official sources inform the method. They do not independently corroborate this project or validate its proposed operating framework.

Angrist, J. D., Imbens, G. W., & Rubin, D. B. (1996). Identification of Causal Effects Using Instrumental Variables. *JASA*, 91, 444–455. Author institutional record. The assumptions for a complier effect exceed random assignment alone. Orientation and altered routing can violate exclusion for a recorded-execution endpoint.

Frangakis, C. E., Rubin, D. B., & Zhou, X. H. (2002). Clustered Encouragement Designs with Individual Noncompliance: Bayesian Inference with Randomization, and Application to Advance Directive Forms. *Biostatistics*, 3, 147–164. Primary abstract. Clustered offers and noncompliance require explicit estimands; this case does not implement the paper's Bayesian estimator.

Moore, G. F., Audrey, S., Barker, M., et al. (2015). Process Evaluation of Complex Interventions: Medical Research Council Guidance. *BMJ*, 350, h1258. Author full text. Investigate implementation, mechanisms and context alongside outcome effects. A delivery measure explains exposure; it cannot replace the customer endpoint.

Carroll, C., Patterson, M., Wood, S., Booth, A., Rick, J., & Balain, S. (2007). A Conceptual Framework for Implementation Fidelity. *Implementation Science*, 2, 40. Primary article. Specification, coverage, dose and delivery quality have distinct roles. The proposed trial retains the assignment population even when delivery is incomplete.

Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20, 37–46. Primary publication. Chance-adjusted coding agreement motivates the proposed taxonomy audit. Agreement, category prevalence and disagreement structure must be reported together; reliability does not establish validity.

NIST/SEMATECH. e-Handbook of Statistical Methods: Individuals Control Charts. Official method. Comparable ordered observations support moving-range limits. The case computes mechanical version-sequence checks privately but does not mistake deliberate schema changes for a stable operating process.

Laney, D. B. (2002). Improved Control Charts for Attributes. *Quality Engineering*, 14, 531–537. Primary publication. Attribute-chart dispersion can require adjustment. This case lacks a suitable repeated count/exposure panel; category totals and rounded percentages cannot supply that missing rate process.

Source scope is stated with each reference. Historical project evidence is the protected register S1; methodological sources establish interpretation and proposed designs, not project results.

RESEARCH REFERENCES / METHODS AND INTERPRETATION

Test the operating mechanism instead of rewarding the metric

Primary scholarly and official sources inform the method. They do not independently corroborate this project or validate its proposed operating framework.

The W. Edwards Deming Institute. The Funnel Experiment. Official explanation of tampering. Reacting to routine variation without understanding the process can worsen results. ARC asks whether an alert supports a feasible intervention; no tampering-induced harm was measured in this project.

Campbell, D. T. (1976). Assessing the Impact of Planned Social Change. Occasional Paper Series, 8. Primary ERIC record. Evaluation must distinguish the social objective from administrative indicators. Catalogue and citation targets can invite optimization of documentation without demonstrating operating improvement.

Ittner, C. D., Larcker, D. F., & Randall, T. (2003). Performance Implications of Strategic Performance Measurement in Financial Services Firms. *Accounting, Organizations and Society*, 28, 715–741. Author institutional record. Measurement-system satisfaction and economic performance differ; simple strategic alignment receives limited support. ARC therefore requires a separate value test.

Tucker, A. L., & Edmondson, A. C. (2003). Why Hospitals Don't Learn from Failures. *California Management Review*, 45(2), 55–72. Author working paper. Immediate workarounds can leave recurring system constraints unresolved. Routing capability problems to accepted owners is a proposed response, not a replication of the study.

OECD/European Union/EC-JRC (2008). Handbook on Constructing Composite Indicators: Methodology and User Guide. Official handbook. Aggregation encodes priorities and requires sensitivity analysis. The source-first ranking here is lexicographic, not a weighted average or calibrated benefit estimate.

Strong, M., Oakley, J. E., Brennan, A., & Breeze, P. (2015). Estimating the Expected Value of Sample Information Using the Probabilistic Sensitivity Analysis Sample. *Medical Decision Making*, 35, 570–583. Author repository. Information value is decision-specific and depends on a model of future evidence. The finite-state example in this case does not use the authors' regression estimator.

Collier, D. (2011). Understanding Process Tracing. *PS: Political Science & Politics*, 44, 823–830. Primary article. Evaluate each contribution hypothesis against diagnostic evidence and alternatives. Repeated summaries cannot make a registry-level claim doubly decisive.

Source scope is stated with each reference. Historical project evidence is the protected register S1; methodological sources establish interpretation and proposed designs, not project results.

The case supports a controlled next test, not an impact claim

[AI] This report records author-led review with parallel evidence, methods and analytical agents. It is not external peer review, client approval or certification.

Major issue	Disposition in the revised case
Source and attribution	Registered instructions support provisional direction and acceptance roles. Original artifacts remain unrecovered; agent-session first person is not assigned to a human actor.
Readiness and scope	Local inclusion, forward labels and deployment are separated. Set bounds replace an invented joint-coverage estimate; persistence-based controls are not dismissed as useless.
Measurement and models	Citation conflicts, untimed cadence, overlapping tags and signed-band order are explicit. Historical scenario arithmetic is not a treatment effect or calibrated probability.
Evaluation and value	The proposed availability-package ITT retains upstream eligibility and mature follow-up. Power and break-even use the same endpoint; variance sensitivity can force redesign.
Decision information	Hypothetical evidence value has explicit states, decisions and likelihoods. It is separated from operating value and does not assume historical customer gains.

Partner critique: can a COO act on the front page?

[AI] Yes: authorize the limited release-validation stage, assign platform and receiving-function owners, and insist on an action-level contract before field evaluation. The new overlap bound changes the interpretation of catalogue coverage without pretending to measure failed adoption.

Minor issues and final publication review

[AI] The review distinguishes nonweekly from daily timing, provenance from authority, signed changes from positive benefit quantiles, and role assignment from accepted capacity. Numerical and rendered-page checks are recorded with the final artifact hashes in the private package.

REVIEW DISPOSITION — PASS WITHIN STATED SCOPE

[AI] Source reconciliation, separate numerical review and publication checks support this registry-based design case. Original artifact contents, historical deployment and customer impact remain unestablished.

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