brillion

Turn noise into decisions.

Enterprise Decision Theory

EDAF™ Volume 2 · Alpha Baseline · Version 0.1.0-alpha · 16 July 2026

Document control

Field	Value
Document ID	EDAF-V02
Title	Enterprise Decision Theory
Version	0.1.0-alpha
Status	Internal alpha - not an approved standard
Publication date	16 July 2026
Framework owner	Brillion (Pty) Ltd
Classification	Technology-independent framework publication

Revision history

Version	Date	Status	Change
0.1.0-alpha	16 July 2026	Internal alpha	Established the initial theory, mathematics, measures and research constructs for controlled review.

STATUS *This publication is an internal alpha baseline. It does not claim external adoption, analyst recognition or status as an international standard. Candidate research constructs require validation before normative use.*

Preface

Enterprise Decision Theory provides the scientific and analytical foundation for the Enterprise Decision Architecture Framework (EDAF™). It defines the structure of a decision problem, distinguishes decision from prediction and action, and establishes the conditions under which probability, utility, loss, optimisation, causal inference and strategic interaction may be used.

The publication is technology independent. It applies to human, rule-based, analytical, optimisation-based, machine-learning and agent-assisted mechanisms when those mechanisms participate in an enterprise decision. DecisionOne EDOS is a reference implementation of EDAF and is outside the normative scope of this volume.

This alpha is deliberately explicit about evidence status. Decision quality, latency, throughput and calibration are established measure families. Decision entropy, decision debt and decision capital are EDAF candidate constructs. They are defined here to enable controlled research; they are not presented as universally accepted scientific measures.

> *An enterprise is fundamentally a governed network of interconnected decisions that transform information into business outcomes.*

Executive summary

A decision exists when an authorised subject selects or determines an alternative for a defined objective under stated information, constraints and uncertainty. A forecast, score or recommendation can inform a decision, but it is not the decision unless it carries the authority and semantics of determination. An action is the execution of an authorised determination. An outcome is the observed effect of that action over a declared horizon.

EDAF uses normative theory to test coherence, descriptive theory to account for real human and organisational behaviour, and prescriptive theory to improve decision practice. No one lens is sufficient. A mathematically coherent policy can be unusable in practice; an observed organisational habit can be common and still violate an objective, constraint or duty.

Decision quality is therefore defined as a governed vector covering objective alignment, logical correctness, evidence sufficiency, authority and control integrity, outcome effectiveness, timeliness, and reconstructability. Outcome alone cannot establish quality because favourable outcomes may follow poor reasoning and unfavourable outcomes may follow sound reasoning under uncertainty.

Enterprise application requires explicit assumptions, model scope, authority, version, temporal horizon, evidence provenance and measures. These requirements connect theory to the canonical Enterprise Decision Object and to later EDAF volumes on architecture, governance, intelligence and DecisionOps.

Table of contents

List of figures

List of tables

Normative language

The key words shall, shall not, should, should not and may express requirement, prohibition, recommendation, discouraged practice and permission respectively. Descriptive statements and candidate constructs do not create conformance requirements unless a clause explicitly uses normative language.

FOUNDATIONAL SCOPE

1. Purpose, scope and conformance

1.1 Purpose

This volume establishes a coherent, auditable theory for enterprise decisions. It provides the conceptual bridge between decision science and the EDAF canonical meta-model. It supplies definitions, mathematical forms, measurement rules and evidence controls that later volumes may reference without redefining core concepts.

1.2 Scope

- individual and collective decisions within an enterprise context;
- decisions under certainty, risk, uncertainty and ambiguity;
- single-stage, sequential and strategically interdependent decisions;
- human, automated and hybrid decision mechanisms;
- ex ante evaluation, runtime measurement and ex post learning.

1.3 Exclusions

This volume does not prescribe a software product, universal objective function, universal risk appetite or sector-specific regulatory interpretation. It does not treat predictive accuracy as sufficient evidence of decision value. It does not claim that every decision can or should be reduced to a scalar utility.

1.4 Conformance

A conformant use of this volume shall identify the decision problem, decision subject, objective, alternatives, material states, information set, consequences, preferences or loss representation, constraints, authority and time horizon. Where a mathematical model is used, assumptions and limitations shall be recorded with the model version.

1.5 Relationship to the EDAF body of knowledge

Table 1 - Volume relationships

EDAF publication	Dependency on this volume	Primary output
Volume 0 - Foundation Standard	Defines canonical terms, principles and conformance controls.	Constitutional baseline
Volume 3 - Enterprise Decision Taxonomy	Uses the decision-problem frame to classify decision types.	Controlled classification
Volume 4 - Enterprise Decision Architecture	Implements theory as decision services, models, records and networks.	Architecture patterns
Volume 5 - Decision Governance	Applies authority, accountability, evidence and risk controls.	Governance system
Volume 6 - Decision Intelligence	Operationalises measures, attribution, simulation and scenario analysis.	Analytical system
Volume 7 - DecisionOps	Controls deployment, experiment, monitoring, drift and improvement.	Operating discipline

PRINCIPLE *EDAF-V02 is theory. It shall remain valid independently of DecisionOne EDOS or any other implementation technology.*

2. Canonical decision problem

2.1 Definition

A decision problem is a governed specification of an objective, decision subject, available alternatives, relevant states, information, consequences, preferences or loss, constraints, authority and temporal horizon. The specification defines what is being determined, for whom or what, under which conditions, and with what accountability.

Table 2 - Canonical decision-problem elements

Element	Required meaning	Control question
Objective	Desired result or value criterion.	What outcome is the decision intended to influence?
Decision subject	Entity, case, resource or condition to which the determination applies.	What is being decided about?
Alternatives	Mutually distinguishable choices available at determination time.	What may be selected or rejected?
State	A relevant condition of the world, observed or unobserved.	Which conditions change consequences?
Information	Evidence and context available before determination.	What was knowable at that time?
Consequences	Effects associated with an alternative and state.	What may follow each choice?
Preferences or loss	Representation of comparative value or harm.	How are consequences compared?
Constraints	Limits that bound feasible alternatives.	What may not be violated?
Authority	Right and accountability to determine or approve.	Who or what may decide?
Horizon	Time over which consequences and outcomes are evaluated.	When is effectiveness assessed?

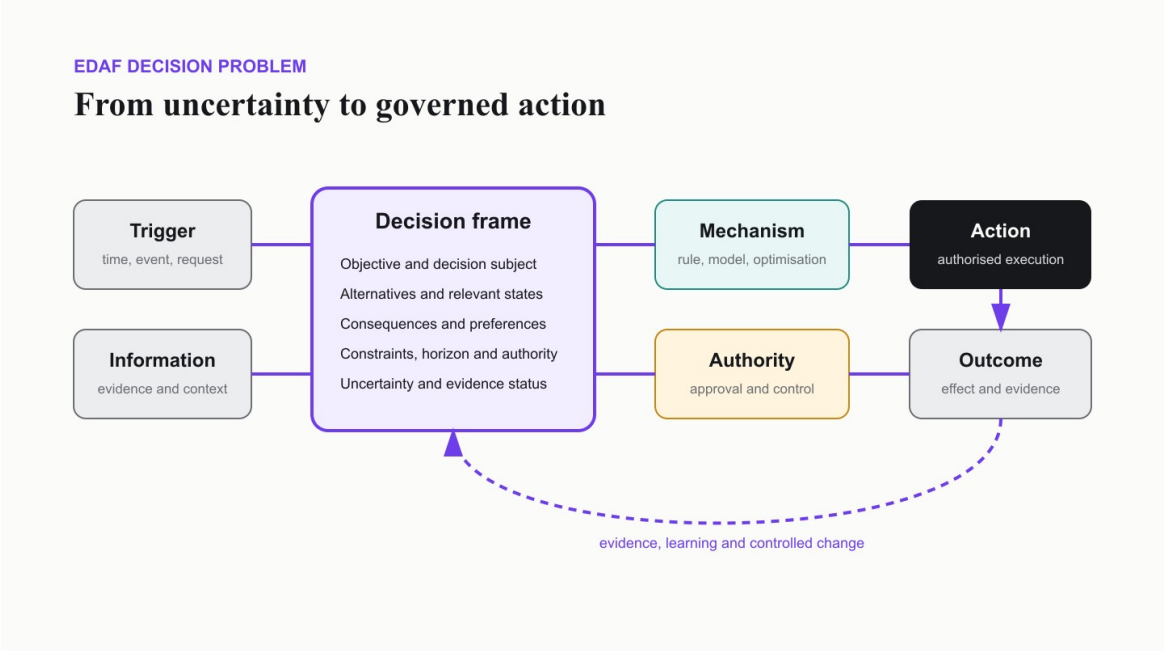

Figure 2 - Canonical decision problem

2.2 Decision, recommendation, action and outcome

A recommendation proposes an alternative. A decision is the authorised determination. An action executes that determination. An outcome is an observed effect. These states may occur in separate systems and at different times. A conformant record shall not merge them when the distinction affects accountability, measurement or audit.

2.3 Prediction is not determination

A predictive model estimates an unknown or future quantity. A decision mechanism uses information, objectives, constraints and authority to determine an alternative. A probability of default, demand forecast or risk score can be an input to several different decisions. Its predictive accuracy does not establish that any resulting policy is economically beneficial, fair, lawful or causally effective.

2.4 Conditions of knowledge

Table 3 - Conditions of knowledge

Condition	Definition	Implication
Certainty	Relevant state and consequence are treated as known.	Choice may be evaluated directly, subject to model validity.
Risk	Relevant outcomes have specified probability distributions.	Expected value or utility may be appropriate.
Uncertainty	Probabilities are incomplete, unstable or not credibly estimable.	Use sensitivity, bounds, scenarios or robust policies.
Ambiguity	The model, state space or interpretation is itself contested.	Record competing models and decision rights for resolution.

2.5 Decision occurrence

The Decision Object is the controlled definition of a decision. A decision instance is one runtime occurrence. A decision record is the evidence created for that occurrence. The definition, instance and record shall carry linked but distinct identifiers. This preserves design-time control while permitting runtime analysis.

> **CAUTION** *A favourable outcome does not prove that a decision was sound. A sound decision may produce an unfavourable outcome when material uncertainty remains.*

3. Decision-theory lenses

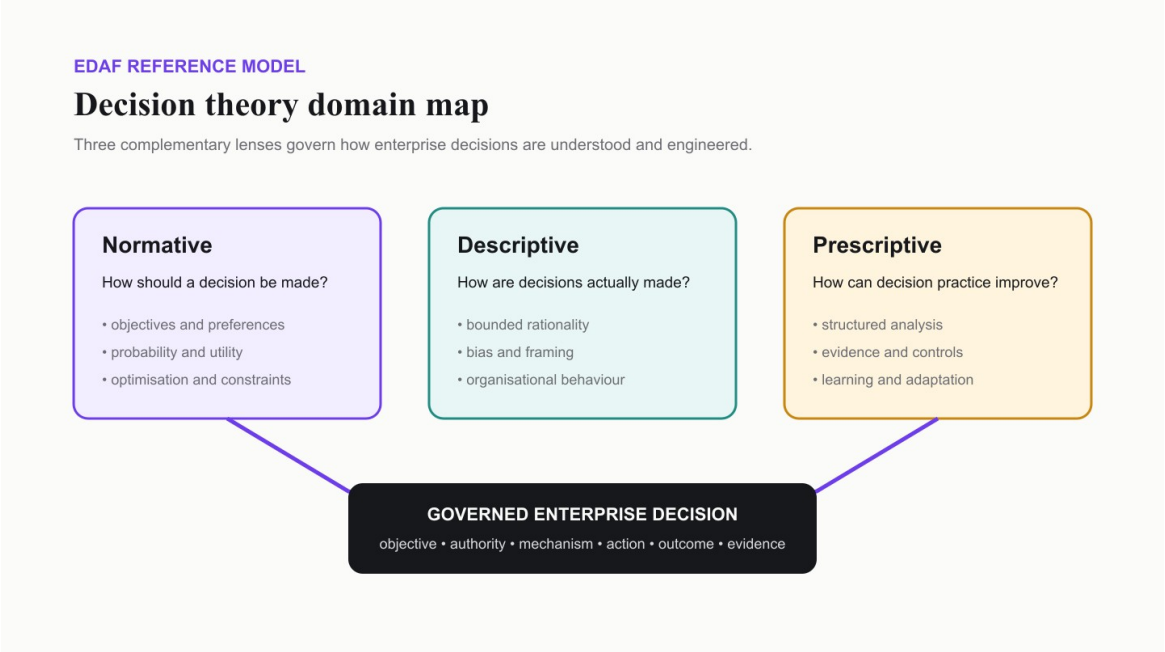

Figure 1 - Decision theory domain map

3.1 Normative theory

Normative theory specifies how a decision should be made under stated axioms and assumptions. Expected utility, Bayesian decision rules and mathematical optimisation are normative forms. Their conclusions are conditional: if objectives, probabilities, preferences and constraints are correctly represented, the prescribed choice follows from the model.

3.2 Descriptive theory

Descriptive theory explains how people and organisations actually decide. Simon’s bounded-rationality model challenges the assumption of complete information and unlimited computation [SRC-SIMON-1955]. Prospect theory describes observed departures from expected-utility behaviour under risk, including reference dependence and asymmetric treatment of gains and losses [SRC-KAHNEMAN-TVERSKY-1979].

These findings inform design and control. They do not, by themselves, define a conformance rule or justify exploiting predictable bias.

3.3 Prescriptive theory

Prescriptive theory develops procedures that improve decisions in practice. Decision analysis structures objectives, alternatives, uncertainty and preferences; operations research supplies mathematical models and computational methods. Arrow described the decision problem in operations research through an objective function, policy alternatives, a model and a computational method [SRC-ARROW-1957].

3.4 EDAF integration rule

Table 4 - Decision-theory lenses

Lens	Primary test	EDAF use	Failure if used alone
Normative	Is the determination coherent under stated assumptions?	Mechanism design and formal verification	May ignore behavioural and organisational limits
Descriptive	What behaviour occurs in context?	Human factors and operating-model design	Common behaviour may remain harmful or incoherent
Prescriptive	What procedure improves practice?	Decision architecture and controlled intervention	Improvement depends on evidence and implementation fidelity

3.5 Rationality and bounded agency

EDAF does not require an enterprise actor to maximise a single complete utility function. It requires the decision architecture to state the value representation actually used, its authority, material exclusions and conflict-resolution method. Where satisficing, lexicographic priorities, thresholds or multi-criteria methods are used, the mechanism shall state those semantics directly.

3.6 Bias, framing and procedural safeguards

- separate evidence from advocacy where material conflicts of interest exist;
- present base rates and uncertainty with point estimates;
- record the option set so omitted alternatives can be audited;
- define escalation when a mechanism encounters out-of-scope conditions;
- test whether labels, ordering or defaults materially change determinations.

INTERPRETATION *Behavioural evidence is a design input. It shall not be converted into a universal catalogue of presumed human defects.*

4. Uncertainty, probability and information

4.1 Probability model

A probability model represents uncertainty over specified events or states. The model shall declare the population, conditioning information, time basis, estimation method and version. Probabilities derived from historical frequencies, expert judgement or model output are not interchangeable unless their semantics and calibration are established.

4.2 Expected value and expected utility

$$\mathrm{EV}(a) = \Sigma_s \, p(s) \, v(c(a,s)) \quad \text{(Eq. 1)}$$

Expected value compares alternatives using the value of consequences. Expected utility replaces direct value with a utility function that represents preferences under risk.

$$\mathrm{EU}(a) = \Sigma_s \, p(s) \, u(c(a,s)) \quad \text{(Eq. 2)}$$

An enterprise shall not infer a shared utility function from monetary amounts alone where outcomes affect different stakeholders, duties or non-fungible constraints.

4.3 Expected loss

$$\mathrm{EL}(a) = \Sigma_s \, p(s) \, L(a,s) \quad \text{(Eq. 3)}$$

A Bayes action minimises expected loss under the specified model. The loss matrix is a governed policy artefact; it is not an objective property discovered by the predictive model.

4.4 Binary cost threshold

For a binary action with zero cost assigned to correct decisions, false-positive cost C_FP and false-negative cost C_FN, acting is preferred when the estimated event probability exceeds the following threshold:

$$p > C\_FP / (C\_FP + C\_FN) \quad \text{(Eq. 4)}$$

This threshold is valid only under the stated assumptions. Action cost, capacity, downstream effects, probability miscalibration, legal restrictions or heterogeneous subjects change the rule.

4.5 Value of information

$$EVPI = E_s[\max_a V(a,s)] - \max_a E_s[V(a,s)] \quad \text{(Eq. 5)}$$

Expected value of perfect information is non-negative under a fixed model because additional information may be ignored. It is an upper bound on what the decision maker should pay for perfect information before considering acquisition cost, delay, implementation limits and model uncertainty.

4.6 Regret

$$R(a,s) = V(a^*(s),s) - V(a,s) \quad \text{(Eq. 6)}$$

Regret compares a selected alternative with the best alternative for the realised state. Minimax regret may be useful when probabilities are not credible, but it expresses a different preference from expected value and shall be identified as such.

4.7 Information entropy

$$H(X) = -\Sigma_x\, p(x) \log_b p(x) \quad \text{(Eq. 7)}$$

Shannon entropy measures uncertainty in a specified probability distribution [SRC-SHANNON-1948]. The logarithm base determines the unit. Entropy is not inherently disorder, risk, decision quality or business value. An EDAF use shall identify the random variable, probability model, logarithm base, population and observation window.

4.8 Information sufficiency

Information is sufficient for a decision only relative to an objective, mechanism and risk tolerance. More data can add cost, delay, bias or privacy exposure without changing the preferred alternative. Information acquisition should therefore be evaluated by expected decision value and control obligations, not by volume alone.

Table 5 - Information sufficiency tests

Information test	Required question	Evidence
Relevance	Can the information change feasible alternatives or their ordering?	Sensitivity or decision-impact analysis
Reliability	Is the source accurate enough for the intended use?	Quality assessment and provenance
Timeliness	Is the information fresh at determination time?	Timestamp, age and validity window
Legitimacy	May the information be used for this purpose?	Policy, consent, legal basis or authority

5. Decision economics

5.1 Economic frame

Decision economics studies how scarce resources, uncertainty, incentives, time and distribution of consequences affect enterprise determinations. A decision architecture shall distinguish accounting cost, economic opportunity cost, expected loss, risk exposure and non-financial constraints.

5.2 Time and discounting

$$PV = \Sigma_t \, CF_t \, / \, (1+r)^t \quad \text{(Eq. 8)}$$

Present value compares cash flows across time using discount rate r. Selection of the rate is a policy decision. It may encode financing cost, opportunity cost and risk assumptions, and it can materially change the ranking of long-horizon alternatives.

5.3 Risk and diversification

Markowitz formalised portfolio selection as a trade-off involving expected return and variance [SRC-MARKOWITZ-1952]. EDAF generalises the architectural lesson: correlated decision exposures shall be evaluated at network or portfolio level where local optimisation can concentrate enterprise risk.

5.4 Incentives and principal-agent effects

A measured target can change the behaviour it measures. Decision rights, performance measures and rewards shall be examined together. A locally rational actor may produce an enterprise-level loss when authority and incentives are misaligned.

5.5 Multi-criteria value

Where consequences cannot credibly be reduced to one monetary amount, the decision model may retain a vector of criteria. Aggregation requires documented scales, weights, vetoes and trade-off rules. A weighted score shall not hide legal, safety or ethical constraints that are non-compensatory.

Table 6 - Value-representation forms

Value form	Appropriate use	Required caution
Monetary value	Consequences are credibly priced and transferable.	Do not omit distribution, externalities or non-financial duties.
Utility	Preferences under risk are explicitly represented.	Validate whose preferences and over which domain.
Loss matrix	Errors or actions have distinguishable costs.	Govern costs independently of predictive training.
Multi-criteria vector	Several objectives remain materially distinct.	Disclose weights, vetoes and sensitivity.
Lexicographic rule	One criterion must dominate another.	Record ordering and exception authority.

5.6 Decision value attribution

Observed business value equals neither model accuracy nor gross outcome. Attribution requires a comparison against a defined counterfactual or baseline, adjustment for implementation cost and an explicit outcome horizon. Where causal identification is not credible, the publication shall describe association rather than causal impact.

5.7 Worked threshold example

Suppose an intervention creates a false-positive cost of 20 units and a false-negative cost of 80 units, while correct classifications are assigned zero incremental cost. Equation 4 yields a threshold of 0.20. The rule is therefore to act when the calibrated event probability exceeds 20 percent.

The result changes if intervention itself costs 10 units, capacity is limited, the benefit varies by subject, or the probability is poorly calibrated. The worked result is an illustration of model semantics; it is not a universal operational threshold.

5.8 Distribution and stakeholder effects

Aggregate benefit can coexist with concentrated harm. A decision assessment should therefore state which stakeholder groups receive costs and benefits, whether effects are reversible, and which obligations constrain trade-offs. Distributional analysis is part of value definition, not a post-processing explanation.

> **CONSTRAINT** *Economic optimisation is subordinate to lawful authority and non-compensatory constraints. A positive expected value does not override a prohibition.*

6. Decision mathematics and mechanism selection

6.1 Constrained optimisation

$$\max_x f(x) \text{ subject to } g_i(x) \leq 0 \text{ and } h_j(x) = 0 \quad \text{(Eq. 9)}$$

The objective function f represents the value being optimised. Inequality and equality constraints define feasibility. A conformant optimisation model shall record the decision variables, units, objective, constraint source, solver status, approximation method and sensitivity to material assumptions.

6.2 Feasibility before optimality

An infeasible recommendation is not improved by a superior objective value. Mechanism design shall test eligibility, authority, resource, timing and safety constraints before ranking feasible alternatives, unless the solver treats those constraints explicitly.

6.3 Mechanism classes

Table 7 - Decision mechanism classes

Mechanism	Primary use	Strength	Control requirement
Policy or rule	Deterministic obligations and thresholds	Transparent semantics	Version, precedence and exception control
Statistical model	Estimate uncertain quantities	Quantified evidence	Population, calibration and drift
Optimisation	Allocate or schedule under constraints	Systematic trade-offs	Objective and feasibility assurance
Simulation	Evaluate scenarios and dynamic effects	Explores uncertainty	Model validation and seed control
Human judgement	Novel, ambiguous or value-laden cases	Contextual interpretation	Competence, authority and record

6.4 Causal decision logic

$$\tau(x) = E[Y(1) - Y(0) \mid X=x] \quad \text{(Eq. 10)}$$

The conditional treatment effect compares potential outcomes under intervention and non-intervention for subjects with characteristics x. Rubin's potential-outcomes framework provides a formal basis for causal effect estimation [SRC-RUBIN-1974]. Because both potential outcomes cannot normally be observed for the same subject, identification depends on design and assumptions.

A model that predicts who will experience an outcome does not necessarily identify who will benefit from an intervention. Decision policies that allocate treatment, retention offers or preventive action should use causal evidence when incremental effect is the relevant objective.

6.5 Robust and stochastic decisions

Stochastic optimisation incorporates probability distributions into the model. Robust optimisation seeks solutions that remain feasible or acceptable across a specified uncertainty set. The selected method shall match the credibility of the available probability model and the enterprise's tolerance for worst-case exposure.

6.6 Mechanism selection test

1. Is the decision repeatable and sufficiently stable to formalise?
2. Are the objective and constraints representable without material distortion?
3. Is there enough evidence to estimate uncertainty or intervention effect?
4. Can the mechanism explain, escalate and fail safely at the required level?
5. Does the expected benefit exceed design, operation and control cost?

6.7 Prediction, optimisation and control

Prediction estimates a quantity. Optimisation selects values for decision variables. Control adjusts action over time in response to observed state. In an enterprise system these may be chained, but each component retains distinct semantics and validation duties.

Table 8 - Analytical component distinctions

Component	Input	Output	Primary validation
Prediction	Features and historical evidence	Estimate or probability	Discrimination, calibration, stability
Decision rule	Evidence, policy and threshold	Determination	Logical correctness and policy conformance
Optimisation	Objective, variables and constraints	Feasible solution	Feasibility, optimality gap, sensitivity
Control	State estimate and feedback	Sequential action	Stability, responsiveness, safety

TRACEABILITY *The final determination shall be traceable through every component that materially changed the alternative set or ordering.*

7. Sequential and strategic decisions

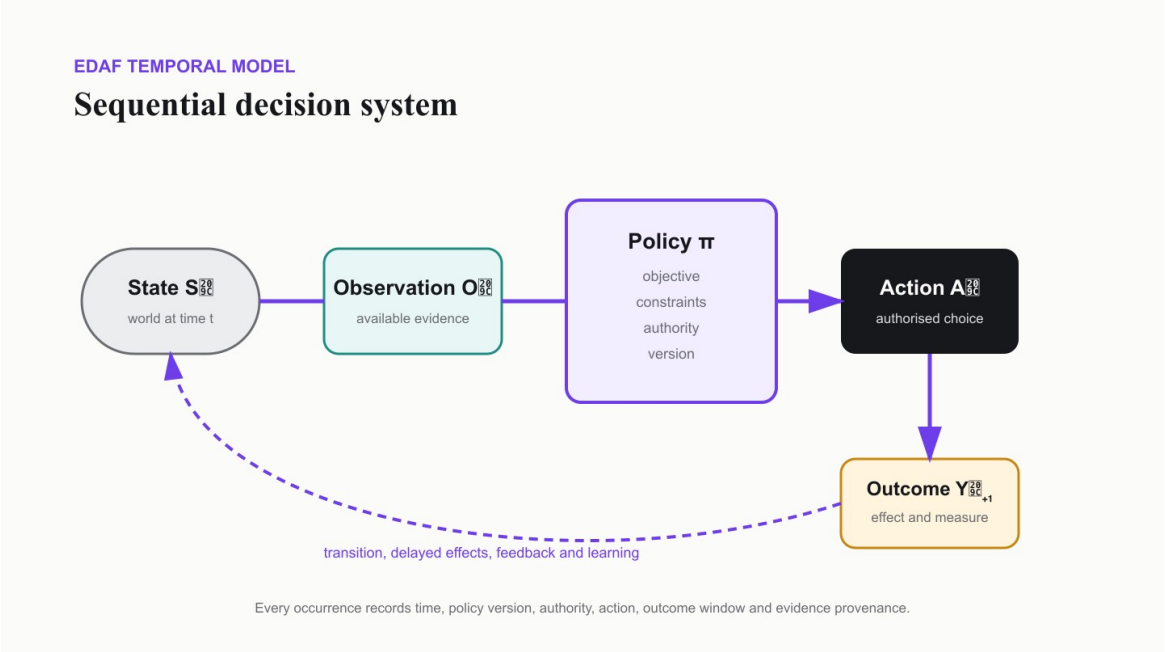

Figure 3 - Sequential decision system

7.1 Sequential model

A sequential decision problem includes state, observation, action, policy, transition and outcome over time. A policy maps available information to an action. The record shall preserve the policy version and information available at each determination time so later evaluation does not introduce hindsight information.

7.2 Delayed effects and temporal credit

An action may create immediate cost and delayed benefit, or shift risk to a later period. Outcome windows shall be declared before evaluation. Where several decisions contribute to one outcome, attribution shall use a stated causal or allocation rule and shall not assign full credit to every contributing decision.

7.3 Exploration and exploitation

A sequential policy may choose between using the best-known alternative and gathering information about uncertain alternatives. Experimentation is governed action. It requires an objective, eligibility criteria, risk bounds, consent or authority where applicable, stopping rules and evidence capture.

7.4 Strategic interaction

Game theory applies when one actor’s outcome depends on the choices of others. Nash proved the existence of equilibrium points for finite games under mixed strategies [SRC-NASH-1950]. An equilibrium is a stability concept; it is not necessarily unique, efficient, fair or desirable.

7.5 Enterprise use

- pricing where customers or competitors adapt;
- fraud controls where adversaries respond to detection policy;
- resource allocation where business units have conflicting incentives;
- negotiation and procurement where information is asymmetric;
- platform policy where participants influence network effects.

7.6 Policy stability and change

A policy may be locally optimal yet unstable when deployed into an adaptive environment. EDAF requires monitoring not only of model inputs and outputs, but also of behavioural responses, constraint pressure and network effects. Policy change shall be versioned and evaluated against a declared baseline.

7.7 Human and automated authority

Table 9 - Authority patterns

Authority pattern	Description	Required control
Human determination	A person makes the authorised choice.	Competence, evidence access and rationale record
Automated determination	A controlled mechanism determines without case-level human approval.	Scope, limits, testing, override and audit
Human approval	A mechanism recommends; a person authorises.	Meaningful review, workload and disagreement capture
Human escalation	Automation determines routine cases and transfers exceptions.	Escalation criteria, service level and safe state

CONTROL INTEGRITY *A nominal human approval step is not meaningful control when the person lacks time, information, authority or a usable means to disagree.*

8. Enterprise decision networks

8.1 Network definition

$$G = (V,E) \quad \text{(Eq. 11)}$$

An Enterprise Decision Network is a directed, typed graph in which nodes V represent governed decision objects and related architectural entities, and edges E represent dependencies such as informs, constrains, authorises, invokes, produces-outcome-for or learns-from.

8.2 Dependency semantics

Table 10 - Decision-network edge types

Edge type	Meaning	Risk if omitted
informs	Evidence from one object enters another decision.	Provenance and propagation cannot be traced.
constrains	One policy or decision limits another's feasible set.	Local decision may violate enterprise control.
authorises	Authority grants determination or execution rights.	Accountability becomes ambiguous.
invokes	One decision triggers another decision service.	Runtime chain cannot be reconstructed.
produces outcome for	An action affects a measured outcome.	Value attribution is disconnected.
learns from	Evidence changes a controlled definition or mechanism.	Change occurs without lineage.

8.3 Network properties

The network view reveals properties that are hidden in isolated decision models: concentration of authority, propagation of stale evidence, circular dependencies, shared constraints, correlated exposures and points where one upstream decision influences many downstream outcomes.

8.4 Local and enterprise objectives

A decision service may optimise a local measure while degrading an enterprise outcome. Conformance therefore requires each material decision objective to trace to a business objective and capability, and requires shared constraints where downstream effects cross organisational boundaries.

8.5 Dependency risk

$$\text{Risk\_path} = P(\text{failure propagates} \mid \text{path}) \times \text{Consequence}(\text{path}) \quad \text{(Eq. 12)}$$

Equation 12 is a generic risk form, not a universal estimator. The probability and consequence shall be defined for the dependency type, time window and control environment. Network risk analysis should identify single points of decision failure and high-centrality nodes whose change has broad consequences.

8.6 Network observability

- stable identifiers for nodes and typed edges;
- version and effective-time semantics;
- correlation IDs across runtime decision chains;
- outcome links and evidence provenance;
- change-impact analysis before controlled release.

9. Decision measurement

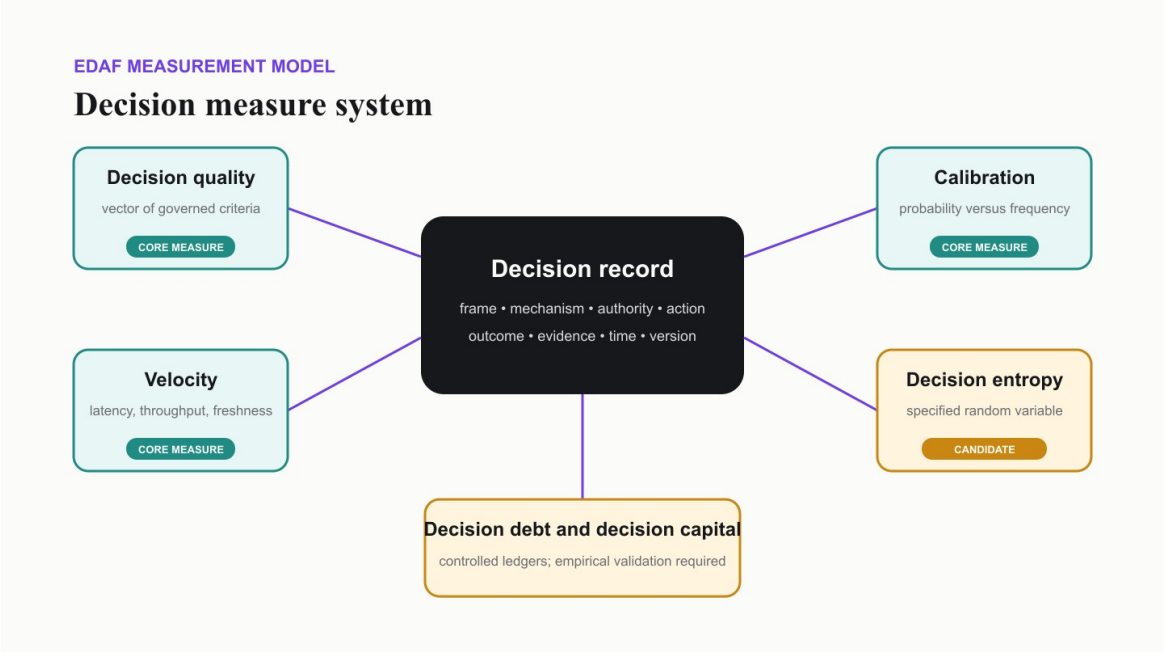

Figure 4 - Decision measure system

9.1 Decision quality vector

$$Q = (q\_obj, q\_logic, q\_evid, q\_auth, q\_out, q\_time, q\_learn) \quad \text{(Eq. 13)}$$

Decision quality is a vector covering objective alignment, logical and constraint correctness, evidence sufficiency, authority and control integrity, outcome effectiveness, timeliness, and learning or reconstructability. Aggregation into one score is permitted only when the component definitions, scales, weights and missing-data rules are controlled.

9.2 Ex ante and ex post quality

Ex ante quality assesses the decision using information available at determination time. Ex post effectiveness uses observed outcomes and a stated attribution method. Both are required. Ex post knowledge shall not be used to rewrite the original information set.

9.3 Decision velocity

$$\text{Latency} = T_{\text{authorised}} - T_{\text{trigger}} \quad \text{(Eq. 14)}$$

$$\text{Throughput} = N_{\text{decisions}} / \Delta t \quad \text{(Eq. 15)}$$

Velocity is a family of temporal measures, including time to determination, time to action, time to outcome, throughput and information freshness. Every reported measure shall state its population, time window, percentile or aggregation method, exclusions and service-level target. Faster is not inherently better when added review reduces material risk.

9.4 Calibration

A probability estimate is calibrated when events assigned a probability occur at approximately that frequency in the relevant population and period. Calibration shall be evaluated over meaningful groups and ranges; aggregate calibration may conceal local error.

$$BS = (1/N) \sum_i (p_i - o_i)^2 \quad \text{(Eq. 16)}$$

The Brier score measures squared error for probabilistic forecasts [SRC-BRIER-1950]. Lower values indicate better performance only within a comparable problem, outcome coding and population.

9.5 Outcome attribution

An outcome measure shall define the observation window, baseline or counterfactual, treatment of delayed effects, competing interventions and censoring. When randomisation is unavailable, the analysis shall state the assumptions used to estimate what would have occurred under a different decision.

9.6 Measurement specification

Table 11 - Decision-measure specification

Field	Requirement
Measure identifier	Stable ID and version
Decision population	Inclusion, exclusion and segmentation rules
Formula	Units, numerator, denominator and aggregation
Time	Trigger, observation window and reporting period
Evidence	Source, lineage, quality and retention
Interpretation	Direction, target, uncertainty and known limitations
Owner	Authority to approve and change the measure

9.7 Measure integrity

A measure that changes incentives may change the process being measured. EDAF therefore requires outcome measures to be reviewed for gaming, selection effects, omitted populations and unintended transfer of cost or risk.

10. EDAF candidate research constructs

> **RESEARCH STATUS** *The constructs in this section are controlled hypotheses for research. They shall not be represented as universally validated measures or used for conformance scoring without an approved validation protocol.*

10.1 Decision entropy

Decision entropy is an EDAF umbrella term for a declared entropy measure applied to a decision-relevant random variable. The term has no valid value without specifying that variable and its probability model.

$$\text{Outcome entropy: } H(Y|I) = -\Sigma_Y\, p(y|I)\, \log_b\, p(y|I) \quad \text{(Eq. 17)}$$

$$\text{Policy dispersion: } H(A|x) = -\Sigma_a\, \pi(a|x)\, \log_b\, \pi(a|x) \quad \text{(Eq. 18)}$$

$$\text{Information gain: } IG = H_{before} - H_{after} \quad \text{(Eq. 19)}$$

Outcome entropy describes residual outcome uncertainty; policy dispersion describes variability in action under a stochastic policy; information gain describes uncertainty reduction. These are not interchangeable. Lower entropy does not necessarily mean a better decision: a consistently harmful policy can have low action entropy.

10.2 Decision debt

Decision debt is a governed ledger of known deficiencies in decision definitions, evidence, mechanisms, controls or records whose remediation has been consciously deferred. It is analogous to an obligation, not a universal monetary balance.

Table 12 - Decision-debt ledger

Ledger field	Meaning	Evidence
Debt item	Specific deficiency or deferred change.	Issue, audit finding or approved exception
Principal estimate	Effort or cost to remediate now.	Estimate with basis and date
Carrying cost	Recurring operating cost caused by deferral.	Incident, rework or manual-control evidence
Risk exposure	Potential consequence while unresolved.	Risk assessment and affected decisions
Due and owner	Accountability and approved horizon.	Authority, review date and disposition

A debt total shall not combine heterogeneous items without controlled scales. Severity, carrying cost and risk exposure should remain separately visible.

10.3 Decision capital

Decision capital is a candidate construct for the governed, reusable capacity that improves the enterprise's ability to make and learn from decisions. Possible components include approved decision definitions, reusable evidence, validated mechanisms, interoperable metadata, competent decision roles and accumulated outcome knowledge.

A future index would require construct validity, non-overlapping components, stable measurement and evidence that the index predicts improved decision capability. Software licences, undocumented code and ungoverned data shall not be counted merely because they exist.

10.4 Candidate relationship model

Decision debt may reduce usable decision capital by consuming maintenance capacity, weakening evidence or constraining reuse. Information gain may reduce uncertainty without reducing debt. These proposed relationships are hypotheses to be tested, not identities.

Table 13 - Candidate construct controls

Construct	Unit or form	Minimum declaration	Prohibited interpretation
Decision entropy	Bits, nats or declared unit	Variable, probability, base, population, window	Universal chaos or quality score
Decision debt	Controlled ledger; optional component estimates	Item, owner, exposure, due date, evidence	Single objective financial liability
Decision capital	Component inventory or validated index	Inclusion, control status, reuse evidence	Count of software or data assets

10.5 Validation gates

1. content validity: the definition covers the intended concept;
2. construct validity: the measure behaves distinctly from adjacent concepts;
3. reliability: repeated measurement is stable under comparable conditions;
4. criterion validity: the construct predicts or explains a relevant external outcome;
5. decision usefulness: use improves choices enough to justify measurement cost.

11. Enterprise application

11.1 Theory-to-object mapping

Table 14 - Theory-to-Decision-Object mapping

Theory element	Canonical Decision Object attribute	Required trace
Objective and value	Business Objective, KPIs	Objective ID and measure specification
Alternatives and mechanism	Decision, Rules, Models, Optimisation	Mechanism version and option set
Information and state	Inputs, Context, Metadata	Evidence provenance and effective time
Constraints and authority	Policies, Constraints, Governance	Policy and authority identifiers
Determination and execution	Recommendation, Actions	Decision instance and action correlation
Consequences and learning	Outcomes, KPIs, Learning, Audit	Outcome horizon and change lineage

11.2 Minimum decision specification

1. decision name, identifier, subject and domain;
2. objective, stakeholder scope and outcome horizon;
3. alternatives, eligibility and constraints;
4. information set, uncertainty and evidence provenance;
5. preference, value or loss semantics;
6. mechanism, authority, escalation and exception path;
7. runtime record, measures, monitoring and learning controls.

11.3 Mechanism assurance

Assurance shall be proportional to decision materiality, automation, reversibility, uncertainty and potential harm. A mechanism can be logically correct and still fail because inputs are stale, authority is invalid, deployment differs from design or outcomes are not observed.

11.4 Human-centred operation

Where people determine, approve or override decisions, architecture shall account for workload, competence, information presentation, time pressure and escalation. Override rate alone does not measure human value; disagreement reason and subsequent outcome are required for interpretation.

11.5 Decision record requirements

Table 15 - Decision-record minimum

Record group	Minimum content
Identity	Decision definition ID, instance ID, subject and correlation ID
Time	Trigger, determination, authorisation, action and outcome timestamps
Evidence	Input references, context, provenance and freshness
Mechanism	Policy, rule, model, optimisation and configuration versions
Authority	Decision right, approver or automated authority and exceptions
Result	Alternatives considered, determination, confidence and explanation
Effect	Action, outcome, KPIs and attribution status

12. Research and validation protocol

12.1 Research questions

- Which decision-quality dimensions predict sustained business outcomes across decision classes?
- When does reduced decision latency improve value, and when does it weaken control?
- Which entropy formulations are stable and useful for particular decision contexts?
- Can decision-debt ledgers predict incident, rework or control cost?
- Can decision-capital components be measured without circularly defining success?

12.2 Study design

Validation should use pre-registered definitions, representative decision populations, declared exclusion rules, temporal holdouts and independent review. Where causal effects are claimed, the design shall identify the counterfactual and threats to identification.

12.3 Evidence classes

Table 16 - EDAF evidence classes

Class	Evidence	Permitted claim
E0	Definition or expert proposal only	Candidate concept
E1	Illustrative cases or internal observations	Feasibility or observed association
E2	Replicated observational study	Stable association within stated scope
E3	Controlled or credible quasi-experimental study	Causal effect within stated assumptions
E4	Independent replication across contexts	Generalised claim within validated domain

12.4 Change control

A construct shall move from candidate to normative only through an approved architecture decision record. The record shall identify evidence class, validated scope, known failure modes, migration impact and the publications that depend on the term.

12.5 Reproducibility

- retain source data or a lawful reproducible extract;
- version code, configuration, formulas and random seeds;
- publish inclusion, exclusion and missing-data rules;
- report uncertainty and negative findings;
- separate exploratory analysis from confirmatory tests.

12.6 Current maturity statement

This alpha establishes terminology, propositions, mathematical forms and controls. It does not provide empirical validation of the candidate EDAF constructs. Subsequent releases shall add domain studies, conformance tests, peer review and cross-industry evidence before any candidate construct is promoted.

> **PUBLICATION CONTROL** *Consistency and traceability take precedence over rapid expansion. A new term shall not enter another EDAF volume without a canonical identifier, definition, evidence status and cross-reference.*

Glossary

Term	Definition
Action	Execution of an authorised determination.
Alternative	A distinguishable choice available to a decision subject.
Calibration	Agreement between stated probabilities and observed frequencies in a defined population and period.
Consequence	A possible effect associated with an alternative and state.
Decision	An authorised determination among alternatives for a defined objective under stated context and constraints.
Decision instance	One runtime occurrence of a defined Decision Object.
Decision Object	The canonical governed representation of an enterprise decision.
Decision policy	A controlled mapping from decision-relevant information to a determination or action.
Decision problem	A specification of objective, subject, alternatives, states, information, consequences, preferences or loss, constraints, authority and horizon.
Decision record	Evidence sufficient to reconstruct and evaluate a decision instance.

Term	Definition
Decision velocity	A family of temporal measures for determination, action, outcome, throughput and information freshness.
Enterprise Decision Network	A directed, typed graph of governed decisions and their dependencies.
Expected utility	Probability-weighted utility of consequences under a specified model.
Information set	Evidence and context available to a decision at determination time.
Loss	A governed representation of harm or cost associated with an alternative and state.
Outcome	An observed effect associated with an action over a declared horizon.
Probability model	A representation of uncertainty over specified events or states.
Recommendation	A proposed alternative that does not itself constitute authorised determination unless explicitly granted that status.
Regret	Difference between the value of the best alternative for a realised state and the selected alternative.
State	A decision-relevant condition of the world, observed or unobserved.
Utility	A representation of preference over consequences under stated conditions.
Value of information	Expected improvement in decision value available from additional information, before acquisition cost.

Appendix A. Notation and equation register

Symbol	Meaning	First use
a, A	Alternative or action; action set where context requires	Eq. 1
s, S	State; state space	Eq. 1
c(a,s)	Consequence of alternative a in state s	Eq. 1
p(s)	Probability assigned to state s	Eq. 1
v, u, L	Value, utility and loss functions	Eq. 1-3
C_FP, C_FN	False-positive and false-negative costs	Eq. 4
H(X)	Entropy of random variable X	Eq. 7
τ(x)	Conditional average treatment effect	Eq. 10
G=(V,E)	Directed graph with node and edge sets	Eq. 11
Q	Decision-quality vector	Eq. 13
BS	Brier score	Eq. 16

Appendix A.1 Equation use

Equations in this volume are reference forms. A conformant implementation shall define units, domains, conditioning information, population, estimation method and version. Symbols may be specialised in later volumes provided their mapping to this register is explicit.

Appendix B. Worked examples

B.1 Capacity-constrained intervention

An enterprise can intervene in at most 1,000 cases. A predictive model estimates event probability, but capacity makes a universal threshold infeasible. The decision problem therefore combines calibrated probability, estimated incremental treatment effect, intervention cost, eligibility policy and a capacity constraint. The resulting ranking is a decision policy, not a prediction.

B.2 Supplier allocation

A procurement decision allocates demand across suppliers. The objective includes expected cost and service performance; constraints include minimum qualification, maximum concentration, capacity and contractual commitments. A lower unit price does not determine the allocation because correlated disruption risk and constraints are enterprise-level effects.

B.3 Human approval

A model recommends a limit increase. A human approver receives the recommendation, evidence, confidence and policy constraints. The decision record distinguishes the recommendation from approval and records disagreement reason. Outcome analysis compares approved and rejected cases using a stated attribution method rather than treating approval as proof of correctness.

Appendix B.4 Scenario comparison

Table 17 - Scenario values

Scenario	Probability	Alternative A value	Alternative B value	Best by state
High demand	0.30	120	170	B
Base demand	0.50	100	95	A
Low demand	0.20	45	10	A

Expected value of A is 95; expected value of B is 100.5. Under this model B has the greater expected value. With perfect information, the expected value is 110.5, so EVPI is 10. The result is conditional on the stated probabilities and values.

Appendix B.5 Quality-versus-outcome matrix

Table 18 - Decision quality and outcome

Ex ante quality	Observed outcome	Interpretation
High	Favourable	Consistent with sound process; attribution still required.
High	Unfavourable	May reflect residual uncertainty or model failure; investigate.
Low	Favourable	Possible good fortune; do not ratify weak process.
Low	Unfavourable	Process and outcome both require remediation.

Appendix C. Evidence register

Source ID	Source	Use	Status
SRC-SIMON-1955	Simon (1955), A Behavioral Model of Rational Choice	Bounded rationality	Verified bibliographic source
SRC-KAHNEMAN-TVERSKY-1979	Kahneman and Tversky (1979), Prospect Theory	Descriptive choice under risk	Verified bibliographic source
SRC-SHANNON-1948	Shannon (1948), A Mathematical Theory of Communication	Information entropy	Verified bibliographic source
SRC-MARKOWITZ-1952	Markowitz (1952), Portfolio Selection	Risk and diversification	Verified bibliographic source
SRC-ARROW-1957	Arrow (1957), Decision Theory and Operations Research	Decision-problem structure	Verified bibliographic source
SRC-NASH-1950	Nash (1950), Equilibrium Points in N-Person Games	Strategic interaction	Verified bibliographic source
SRC-RUBIN-1974	Rubin (1974), Estimating Causal Effects	Potential outcomes	Verified bibliographic source
SRC-BRIER-1950	Brier (1950), Verification of Forecasts	Probability-score measurement	Verified bibliographic source

Bibliography

Arrow, K. J. (1957). Decision Theory and Operations Research. Operations Research, 5(6), 765-774. https://doi.org/10.1287/opre.5.6.765

Brier, G. W. (1950). Verification of Forecasts Expressed in Terms of Probability. Monthly Weather Review, 78(1), 1-3. https://doi.org/10.1175/1520-0493(1950)078<0001:VOFEIT>2.0.CO;2

Howard, R. A. (1966). Decision Analysis: Applied Decision Theory. Stanford Research Institute.

INFORMS. Decision Analysis. Official society subject-area description. Accessed 16 July 2026.

Kahneman, D., and Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. Econometrica, 47(2), 263-291. https://doi.org/10.2307/1914185

Markowitz, H. (1952). Portfolio Selection. Journal of Finance, 7(1), 77-91. https://doi.org/10.1111/j.1540-6261.1952.tb01525.x

Nash, J. F. (1950). Equilibrium Points in N-Person Games. Proceedings of the National Academy of Sciences, 36(1), 48-49. https://doi.org/10.1073/pnas.36.1.48

Rubin, D. B. (1974). Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies. Journal of Educational Psychology, 66(5), 688-701. https://doi.org/10.1037/h0037350

Shannon, C. E. (1948). A Mathematical Theory of Communication. Bell System Technical Journal, 27(3), 379-423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

Simon, H. A. (1955). A Behavioral Model of Rational Choice. Quarterly Journal of Economics, 69(1), 99-118. https://doi.org/10.2307/1884852

Index

brillion

Turn noise into decisions.

Brillion (Pty) Ltd · brillion.africa